Appendix C

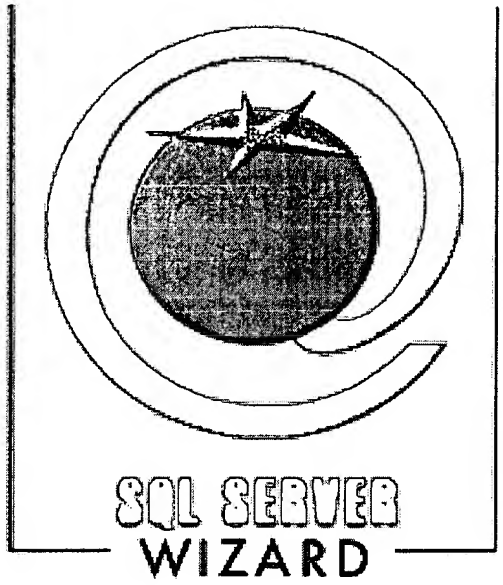

PROLIANT

Introduction

Summary

Configuration Data

CPU Consumption

M mory

Disk Occupation

Disk I/O

Connection Count rs

Hit Rate

Logs

Backup

Concepts

Introduction

Based on data collected in the host proliant, from 08/10/2001, at 00:00, to 08/18/2001, at 23:00, the current performance analysis report was elaborated for SQL Server instance PROLIANT.

The data used in this report was obtained from an exclusive collector, developed specially for this end, executing on the target instance with high resolution and low intrusion. This collector obtains data directly from the SQL Server instance, without any other libraries or additional tools, with a minimum overhead on the system. The data collected is stored using a binary format, in order to provide persistence. When automatically sent, it is compressed and encrypted, to ensure fast delivery and confidentiality.

The content of this report is based on years of experience in performance analysis and capacity planning. The tool used to generate this report operates in a completely automatic way, without direct human intervention. It uses an extensible inference machine, based on heuristics and rules, and is subject to continuous improvements. Using concepts such as "watermarks" and tolerance, it is possible to determine if a computational resource usage is excessive and if the excess is relevant.

During the monitoring period, the summary configuration of the instance, which has been obtained dynamically, was:

Name : Microsoft SQL Server 2000
Version : 8.00.194 (Intel X86)
Edition : Enterprise Evaluation Edition

Summary

This report refers to the period from 08/10/2001, at 00:00, to 08/18/2001, at 23:00. The following highlights were registered:

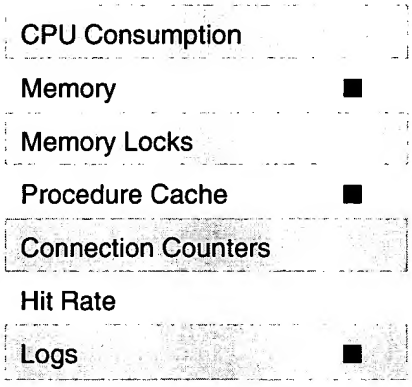

The procedure cache hit rate remained low for most of the monitored period.

The log usage rate remained high throughout the monitored period.

Configuration Data

The table below shows the main configuration parameters for SQL Server instance PROLIANT. The columns indicate the name of the parameter, its status, configured value, current value, minimum and maximum values.

Parameter	Status	Configured Value	Current Value	Minimum Value	Maximum Value
affinity mask	Static and Advanced	0	0	0	2147483647
allow updates	Dynamic	0	0	0	1
awe enabled	Static and Advanced	0	0	0	1
c2 audit mode	Static and Advanced	0	0	0	1
cost threshold for parallelism	Dynamic and Advanced	5	5	0	32767
cursor threshold	Dynamic and Advanced	-1	-1	-1	2147483647
default full-text language	Dynamic and Advanced	1033	1033	0	2147483647
default language	Dynamic	0	0	0	9999
fill factor (%)	Static and Advanced	0	0	0	100
index create memory (KB)	Dynamic and Advanced	0	0	704	2147483647
lightweight pooling	Static and Advanced	0	0	0	1
locks	Static and Advanced	0	0	5000	2147483647
max degree of parallelism	Dynamic and Advanced	0	0	0	32
max server memory (MB)	Dynamic and Advanced	131	131	4	2147483647
max text repl size (B)	Dynamic	65536	65536	0	2147483647
max worker threads	Static and Advanced	255	255	32	32767

Configuration Data

Parameter	Status	Configured Value	Current Value	Minimum Value	Maximum Value
media retention	Static and Advanced	0	0	0	365
min memory per query (KB)	Dynamic and Advanced	1024	1024	512	2147483647
min server memory (MB)	Dynamic and Advanced	32	32	0	2147483647
nested triggers	Dynamic	1	1	0	1
network packet size (B)	Dynamic and Advanced	4096	4096	512	65536
open objects	Static and Advanced	0	0	0	2147483647
priority boost	Static and Advanced	1	1	0	1
query governor cost limit	Dynamic and Advanced	0	0	0	2147483647
query wait (s)	Dynamic and Advanced	-1	-1	-1	2147483647
recovery interval (min)	Dynamic and Advanced	0	0	0	32767
remote access	Static	1	1	0	1
remote login timeout (s)	Dynamic	20	20	0	2147483647
remote proc trans	Dynamic	0	0	0	1
remote query timeout (s)	Dynamic	600	600	0	2147483647
scan for startup procs	Static and Advanced	0	0	0	1
set working set size	Static and Advanced	0	0	0	1
show advanced options	Dynamic	1	1	0	1
two digit year cutoff	Dynamic and Advanced	2049	2049	1753	9999
user connections	Static and Advanced	0	0	0	32767

Configuration Data

Param ter	Status	Configured Value	Current Value	Minimum Value	Maximum Valu
user options	Dynamic	0	0	0	32767

CPU Consumption

CPU usage was low all the time, and SQL Server was responsible, at the most, for 5.9%.

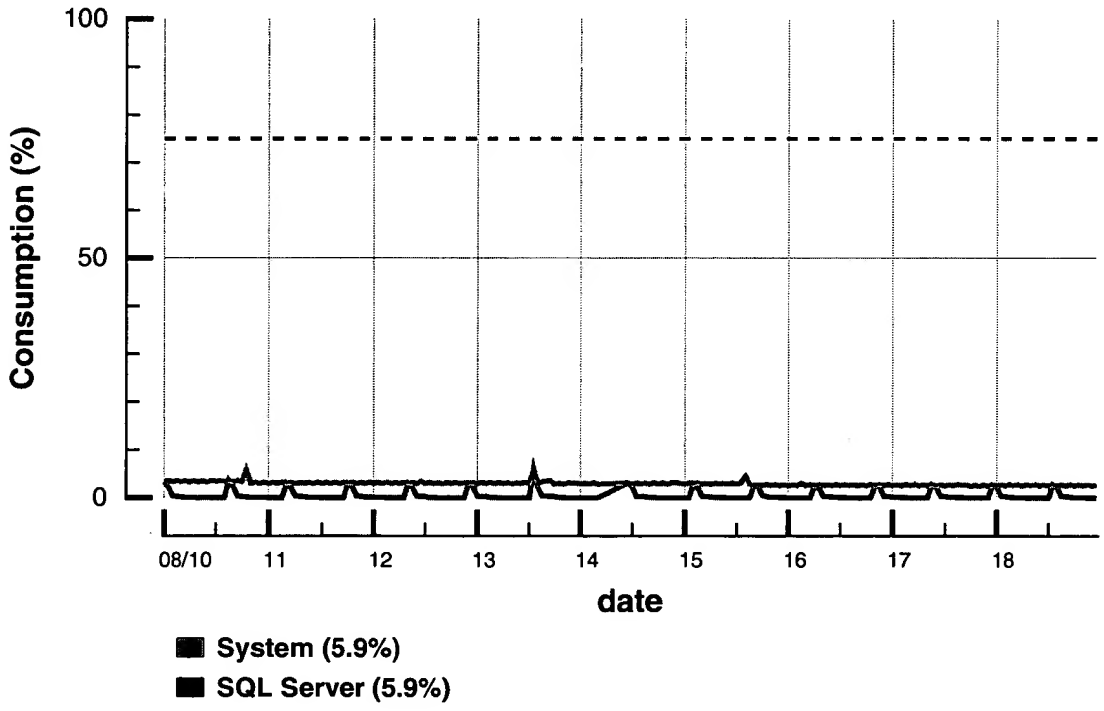

Memory

The server's memory consumption was low most of the time, so it was not a problem.

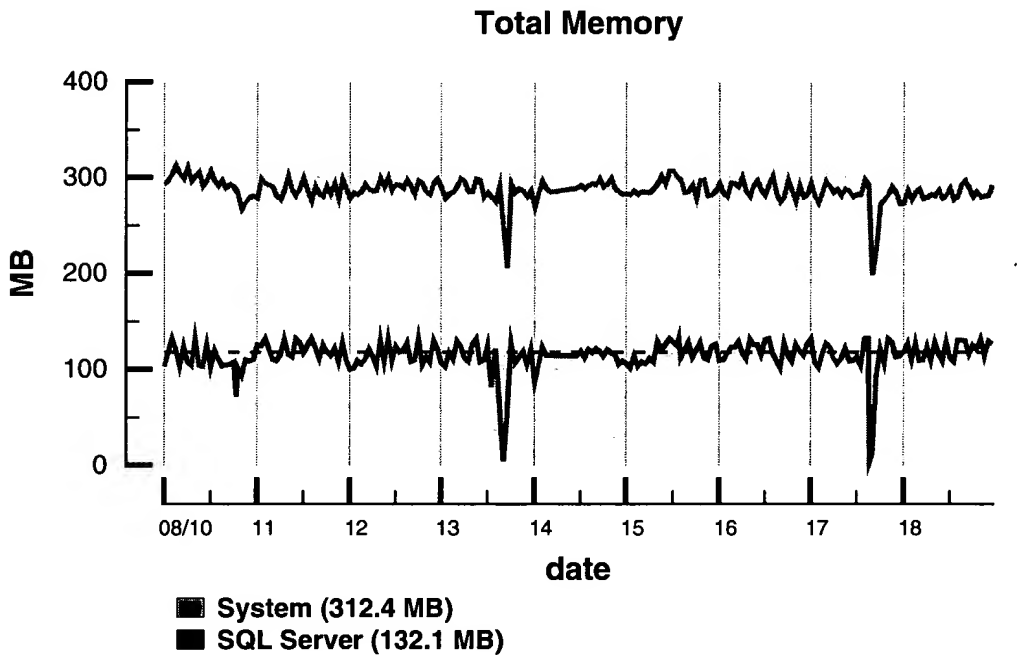

Locked memory remained low throughout the monitored period, so this was not a problem.

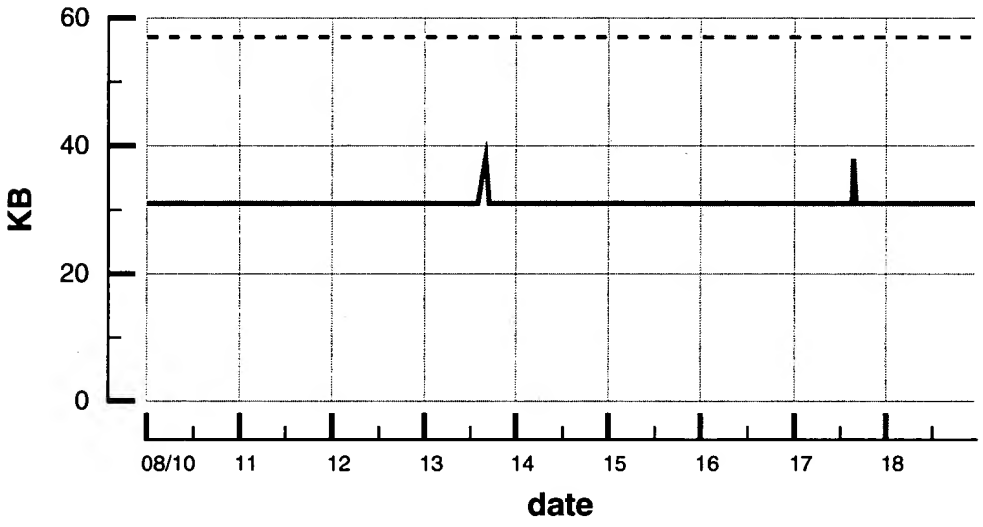

The Procedure Cache had a low memory usage throughout the period, but it exceeded the thershold on occasions, and this may have caused constraint.

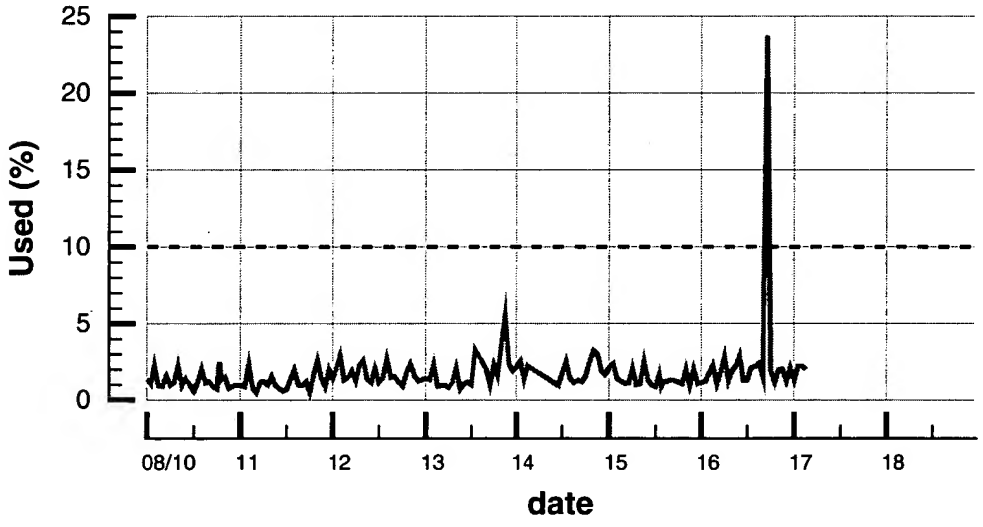

Disk Occupation

The data below refers to the SQL Server data files' disk utilization at the end of the monitoring period. All disks are shown, with the free and the used percentages.

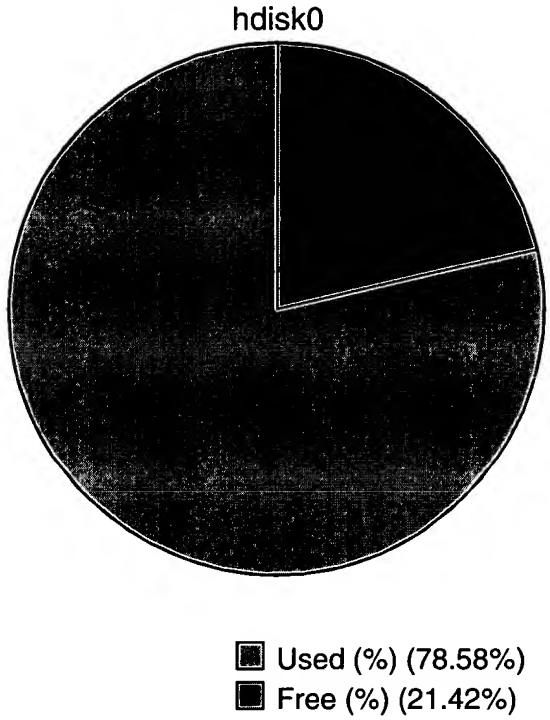

Disk Occupation

Below is a list of all databases for SQL Server instance PROLIANT, along with the database name, size and the list of log and data files.

Database	Name	File name	Size (KB)
copiaDB2	copiaDB2_dat	C:\Program Files\Microsoft SQL Server\MSSQL\data\copiaDB2.mdf	307,200.00
	copiaDB2_log	C:\Program Files\Microsoft SQL Server\MSSQL\data\copiaDB2.ldf	102,400.00
insight_db_V2	insight_db_V2	C:\Program Files\Microsoft SQL Server\MSSQL\data\insight_db_V2.mdf	30,720.00
	insight_db_V2Log	C:\Program Files\Microsoft SQL Server\MSSQL\data\insight_db_V2Log.ldf	51,080.00
master	master	C:\Program Files\Microsoft SQL Server\MSSQL\data\master.mdf	108,992.00
	mastlog	C:\Program Files\Microsoft SQL Server\MSSQL\data\mastlog.ldf	33,088.00
model	modeldev	C:\Program Files\Microsoft SQL Server\MSSQL\data\model.mdf	640.00
	modellog	C:\Program Files\Microsoft SQL Server\MSSQL\data\modellog.ldf	512.00
msdb	MSDBData	C:\Program Files\Microsoft SQL Server\MSSQL\data\msdbdata.mdf	25,280.00
	MSDBLog	C:\Program Files\Microsoft SQL Server\MSSQL\data\msdblog.ldf	10,752.00
newdb2cop	newdb2cop_Data	D:\MSSQL7\Data\newdb2cop_Data.MDF	30,720.00
	newdb2cop_Log	D:\MSSQL7\Data\newdb2cop_Log.LDF	40,712.00
NEWDB2	NEWDB2_1_Data	C:\MSSQL7\Data\NewDB2c	256,000.00
	NEWDB2_1_Log	C:\MSSQL7\Data\newdb2b_log	327,688.00
	NEWDB2_Data	D:\MSSQL7\Data\NEWDB2_Data.MDF	409,600.00

Disk Occupation

Database	Name	File name	Siz (KB)
NEWDB2	NEWDB2_Log	D:\MSSQL7\Data\NEWDB2_Log.LDF	65,736.00
Northwind	Northwind	C:\Program Files\Microsoft SQL Server\MSSQL\data\northwnd.mdf	3,008.00
	Northwind_log	C:\Program Files\Microsoft SQL Server\MSSQL\data\northwnd.ldf	2,048.00
NWCOPY	nwcopy_Data	C:\MSSQL7\data\nwcopy_Data.MDF	1,792.00
	nwcopy_Log	C:\MSSQL7\data\nwcopy_Log.LDF	1,024.00
	NWCOPY_1_Data	D:\MSSQL7\Data\nwcopy2_Data.MDF	20,480.00
pubs	pubs	C:\Program Files\Microsoft SQL Server\MSSQL\data\pubs.mdf	1,408.00
	pubs_log	C:\Program Files\Microsoft SQL Server\MSSQL\data\pubs_log.ldf	768.00
tempdb	tempdev	C:\Program Files\Microsoft SQL Server\MSSQL\data\tempdb.mdf	102,400.00
	templog	C:\Program Files\Microsoft SQL Server\MSSQL\data\templog.ldf	51,200.00

Disk I/O

The table below shows the I/O data for the disks used by the SQL Server, along with the data files in each disk and their locations.

Disk	KB/sec	Data file	Location
hdisk0	146.08	copiaDB2_dat	C:\Program Files\Microsoft SQL Server\MSSQL\data\copiaDB2.mdf
hdisk0	146.08	copiaDB2_log	C:\Program Files\Microsoft SQL Server\MSSQL\data\copiaDB2.ldf
hdisk0	146.08	insight_db_V2	C:\Program Files\Microsoft SQL Server\MSSQL\data\insight_db_V2.mdf
hdisk0	146.08	insight_db_V2Log	C:\Program Files\Microsoft SQL Server\MSSQL\data\insight_db_V2Log.ldf
hdisk0	146.08	master	C:\Program Files\Microsoft SQL Server\MSSQL\data\master.mdf
hdisk0	146.08	mastlog	C:\Program Files\Microsoft SQL Server\MSSQL\data\mastlog.ldf
hdisk0	146.08	modeldev	C:\Program Files\Microsoft SQL Server\MSSQL\data\model.mdf
hdisk0	146.08	modellog	C:\Program Files\Microsoft SQL Server\MSSQL\data\modellog.ldf
hdisk0	146.08	MSDBData	C:\Program Files\Microsoft SQL Server\MSSQL\data\msdbdata.mdf
hdisk0	146.08	MSDBLog	C:\Program Files\Microsoft SQL Server\MSSQL\data\msdblog.ldf
hdisk0	146.08	newdb2cop_Data	D:\MSSQL7\Data\newdb2cop_Data.MDF
hdisk0	146.08	newdb2cop_Log	D:\MSSQL7\Data\newdb2cop_Log.LDF
hdisk0	146.08	NEWDB2_1_Data	C:\MSSQL7\Data\NewDB2c
hdisk0	146.08	NEWDB2_1_Log	C:\MSSQL7\Data\newdb2b_log
hdisk0	146.08	NEWDB2_Data	D:\MSSQL7\Data\NEWDB2_Data.MDF
hdisk0	146.08	NEWDB2_Log	D:\MSSQL7\Data\NEWDB2_Log.LDF

Disk I/O

Disk	KB/sec	Data file	Location
hdisk0	146.08	Northwind	C:\Program Files\Microsoft SQL Server\MSSQL\data\northwnd.mdf
hdisk0	146.08	Northwind_log	C:\Program Files\Microsoft SQL Server\MSSQL\data\northwnd.ldf
hdisk0	146.08	nwcopy_Data	C:\MSSQL7\data\nwcopy_Data.MDF
hdisk0	146.08	nwcopy_Log	C:\MSSQL7\data\nwcopy_Log.LDF
hdisk0	146.08	NWCOPY_1_Data	D:\MSSQL7\Data\nwcopy2_Data.MDF
hdisk0	146.08	pubs	C:\Program Files\Microsoft SQL Server\MSSQL\data\pubs.mdf
hdisk0	146.08	pubs_log	C:\Program Files\Microsoft SQL Server\MSSQL\data\pubs_log.ldf
hdisk0	146.08	tempdev	C:\Program Files\Microsoft SQL Server\MSSQL\data\tempdb.mdf
hdisk0	146.08	templog	C:\Program Files\Microsoft SQL Server\MSSQL\data\templog.ldf

Below is the list of disks that contain at least one SQL Server data file, along with how much I/O each one is responsible for, as a percentage of the total I/O.

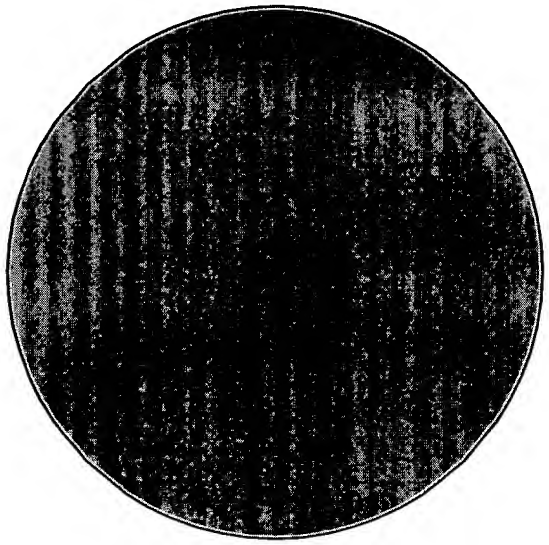

hdisk0 (100.00%)

Connection Counters

The graph below shows the number of connections made to the SQL Server during the monitored period. The limit of simultaneous connections to the database is "unlimited".

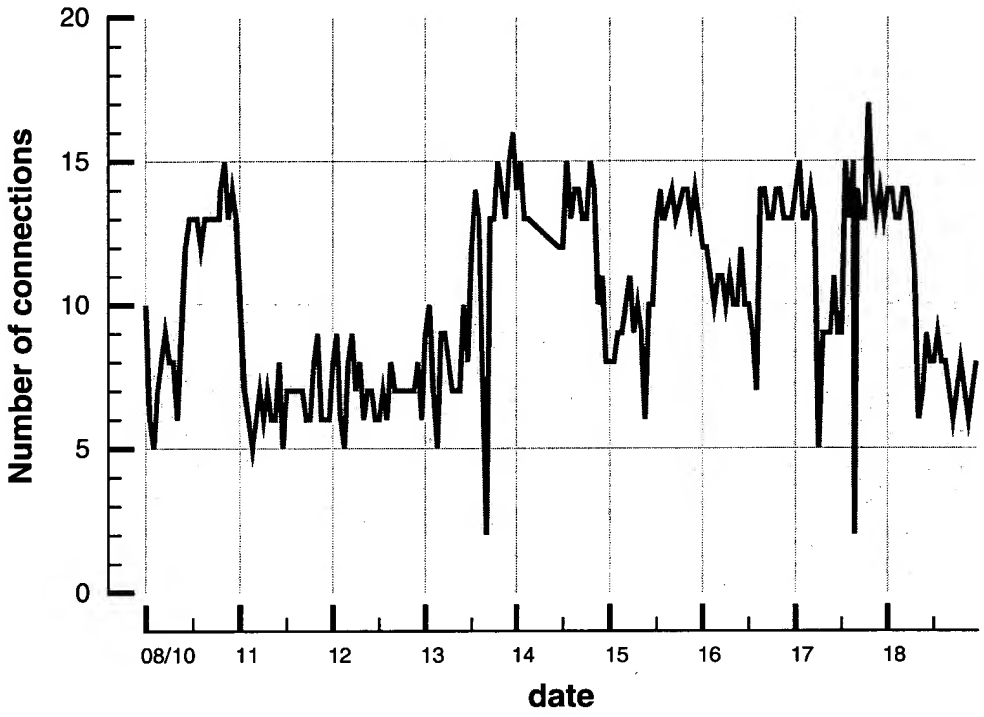

The buffer cache hit rate was high all the time, indicating that the server's memory is sufficient.

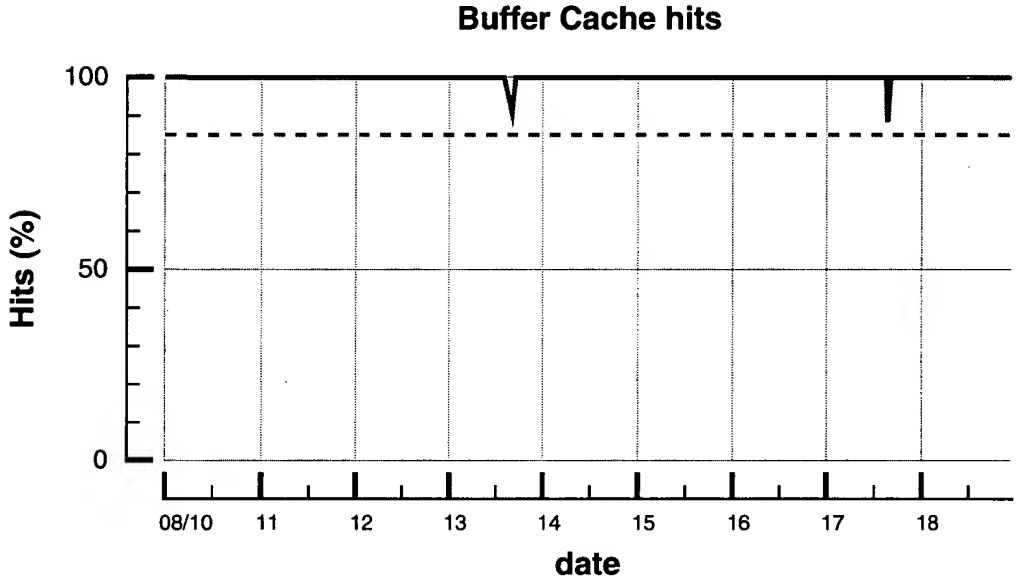

The log cache hit rate remained low throughout the monitored period, which probably indicates a memory shortage for the SQL Server.

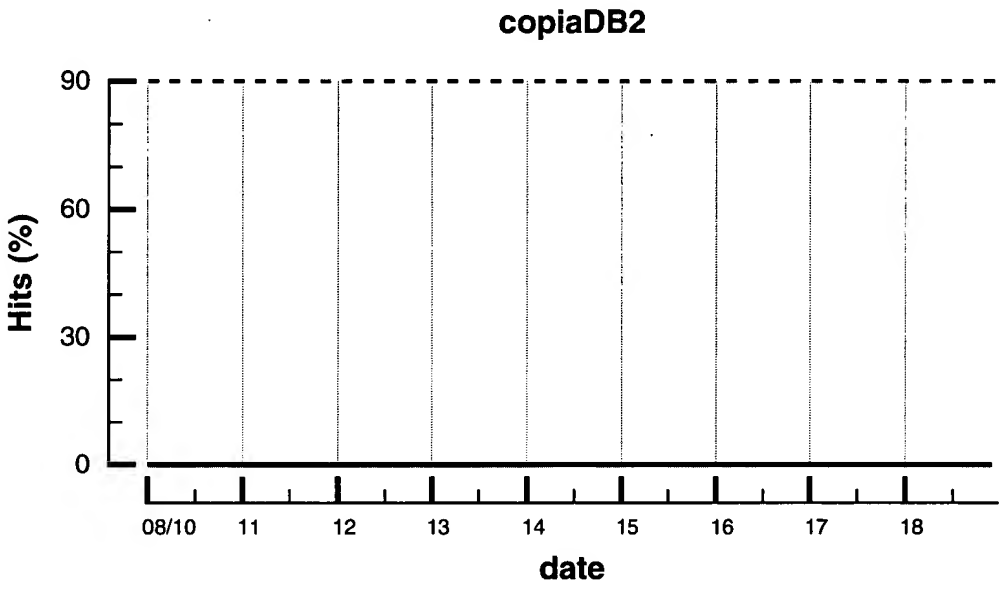

The log cache hit rate remained low throughout the monitored period, which probably indicates a memory shortage for the SQL Server.

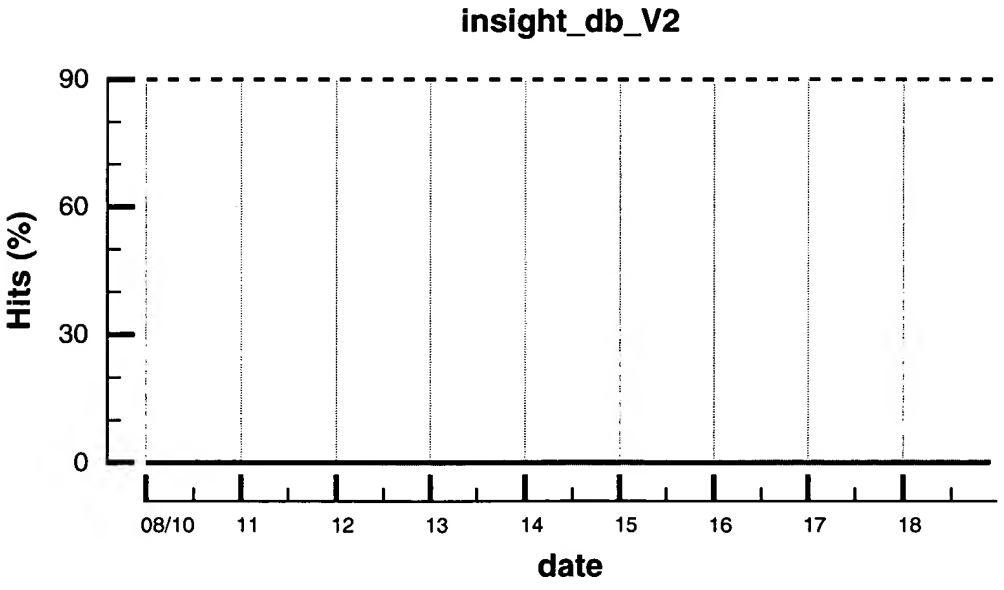

The log cache hit rate remained low throughout the monitored period, which probably indicates a memory shortage for the SQL Server.

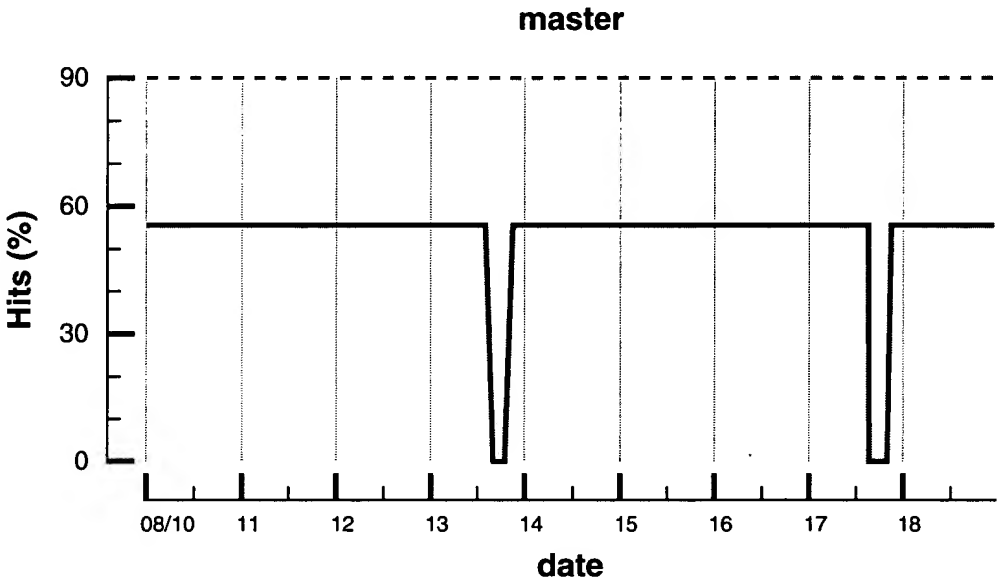

The log cache hit rate remained low throughout the monitored period, which probably indicates a memory shortage for the SQL Server.

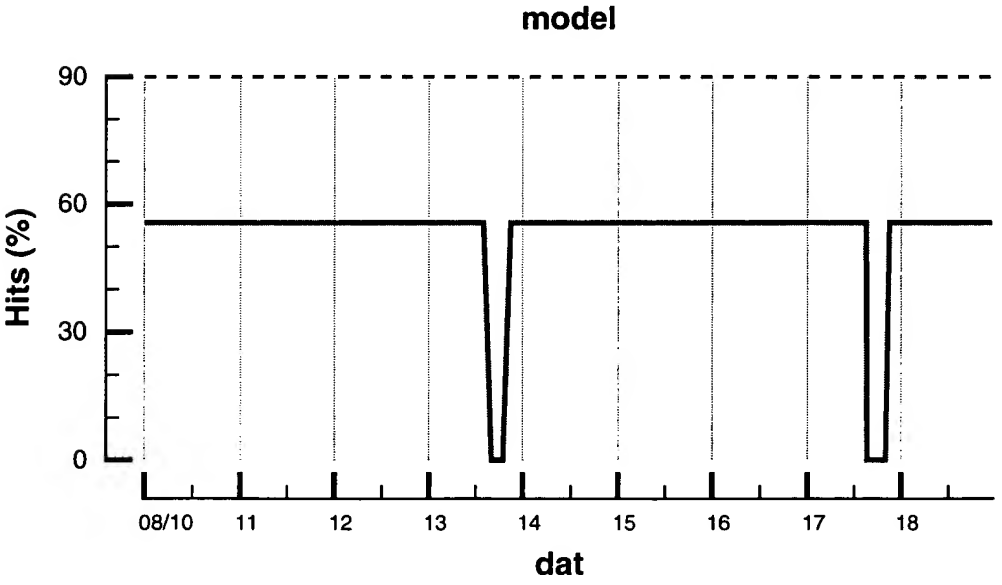

The log cache hit rate remained low throughout the monitored period, which probably indicates a memory shortage for the SQL Server.

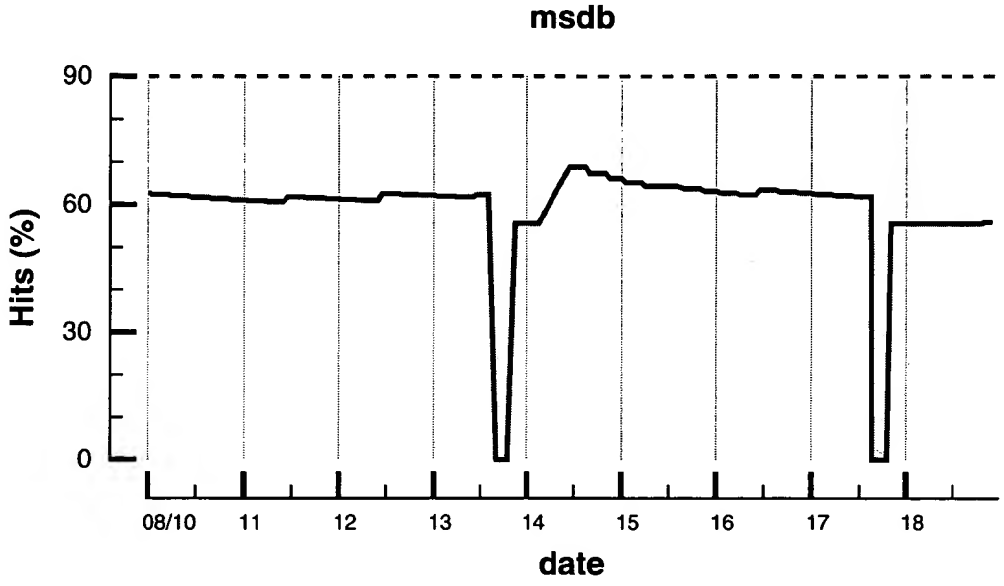

The log cache hit rate remained low throughout the monitored period, which probably indicates a memory shortage for the SQL Server.

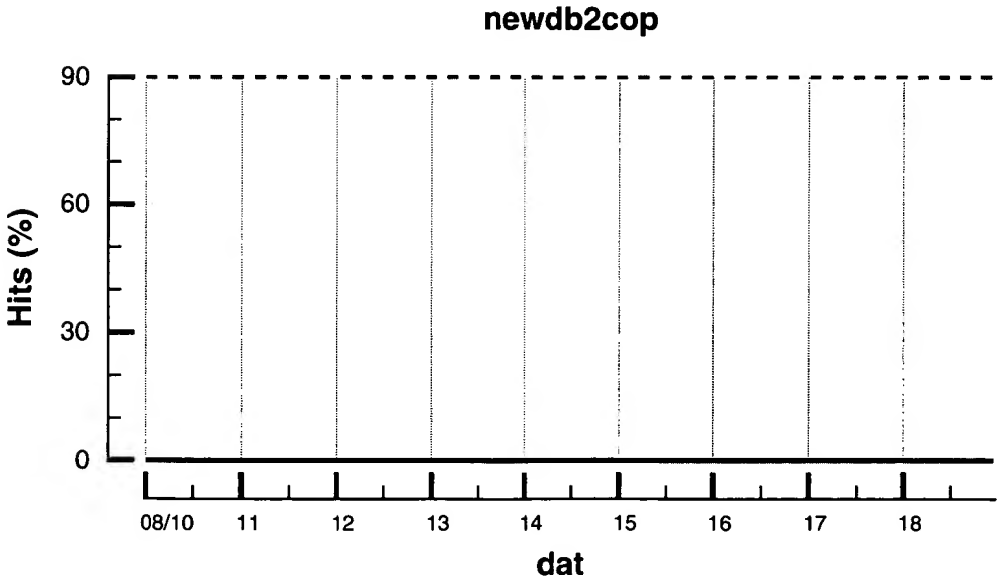

The log cache hit rate remained low throughout the monitored period, which probably indicates a memory shortage for the SQL Server.

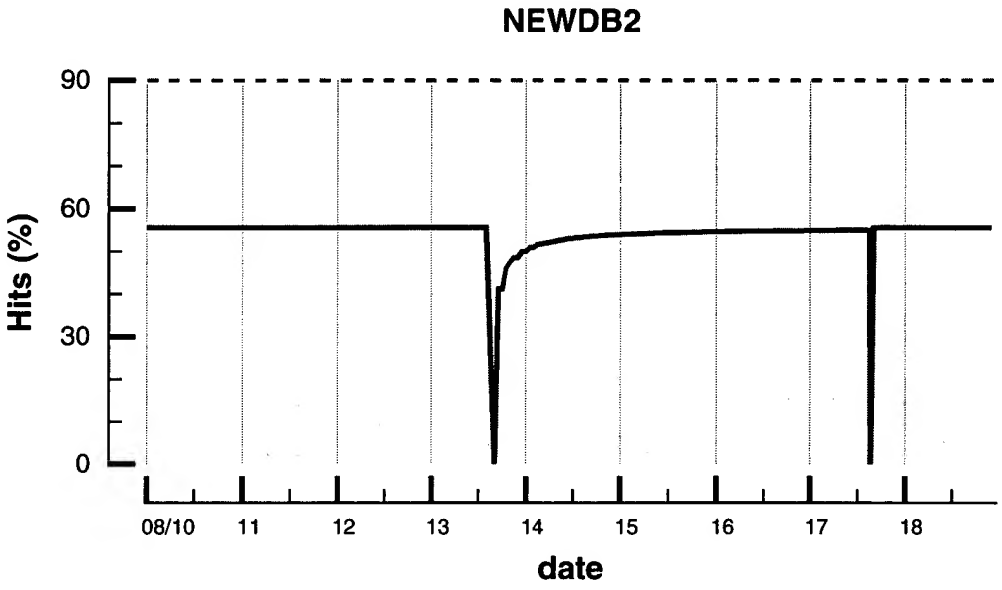

The log cache hit rate remained low throughout the monitored period, which probably indicates a memory shortage for the SQL Server.

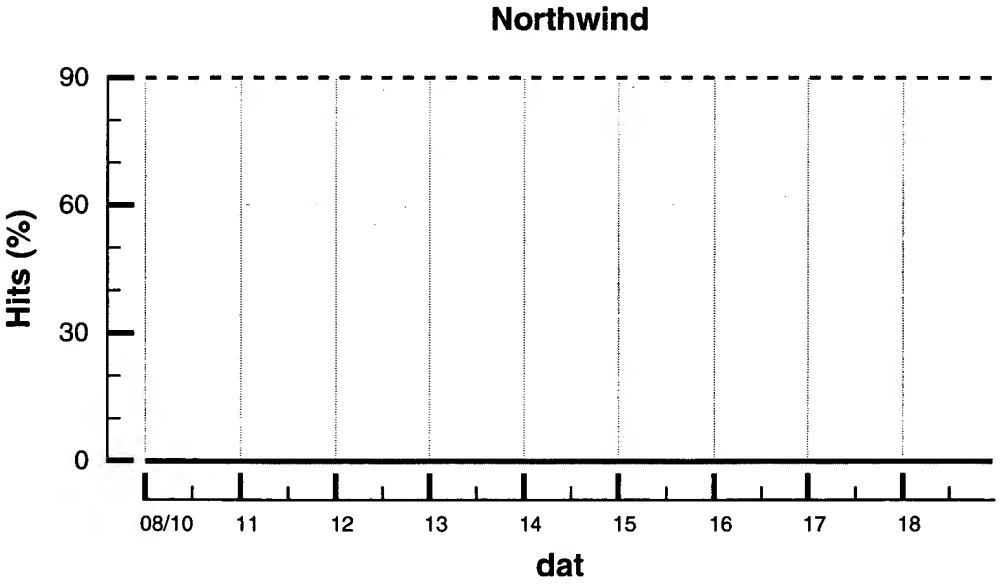

The log cache hit rate remained low throughout the monitored period, which probably indicates a memory shortage for the SQL Server.

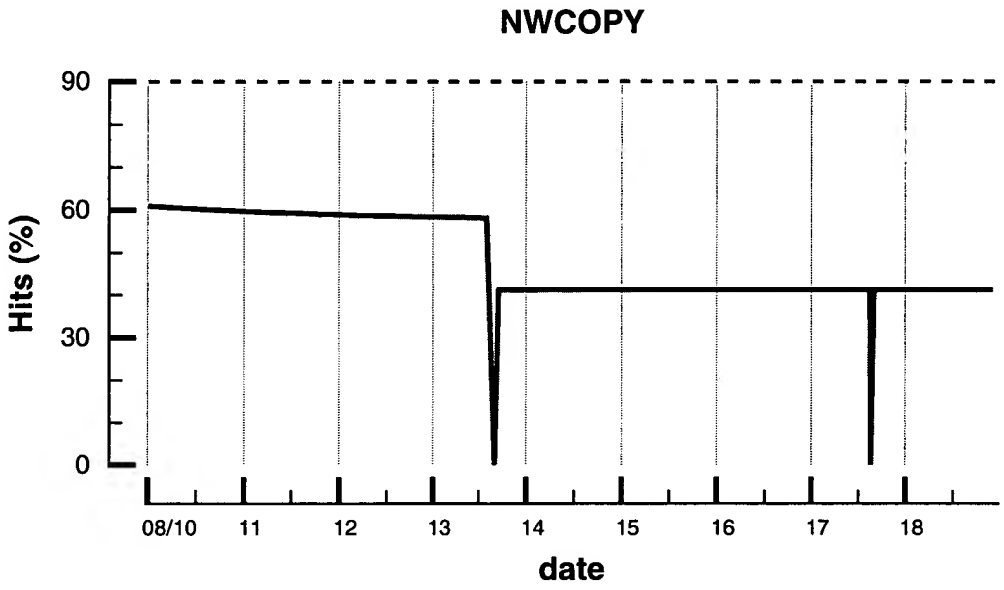

The log cache hit rate remained low throughout the monitored period, which probably indicates a memory shortage for the SQL Server.

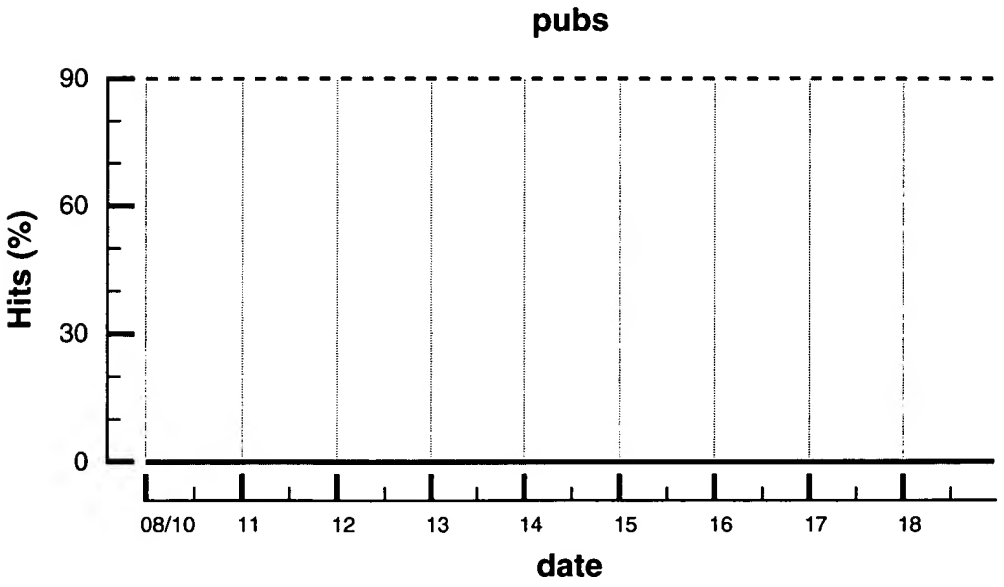

The log cache hit rate remained low throughout the monitored period, which probably indicates a memory shortage for the SQL Server.

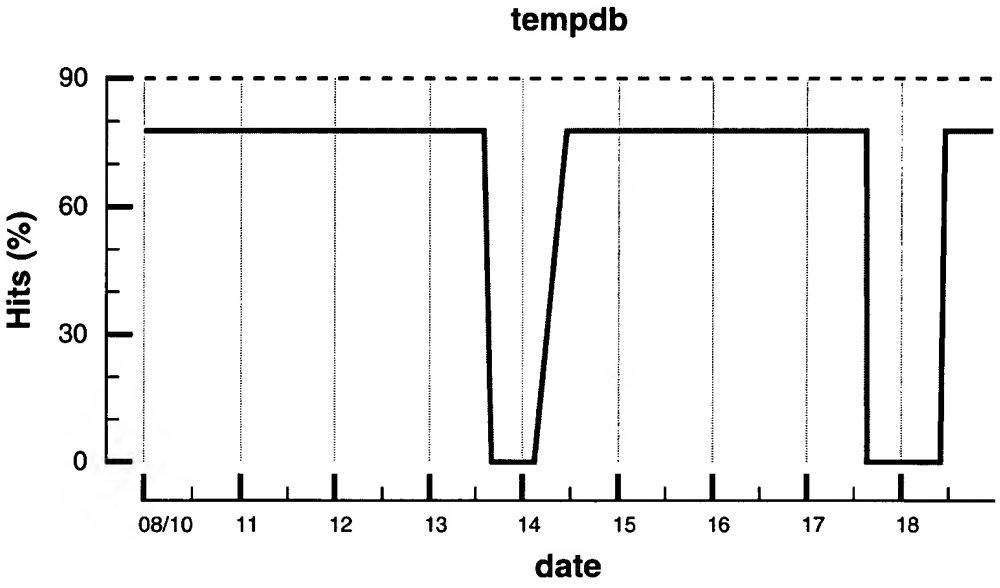

Below are the database log occupation rates. For each database, there is a graph indicating the evolution of these logs during the monitored period, and another graph showing the maximum daily occupation. Only the last 7 monitored days are shown

The log usage remained low during the whole monitored period, not indicating any problems.

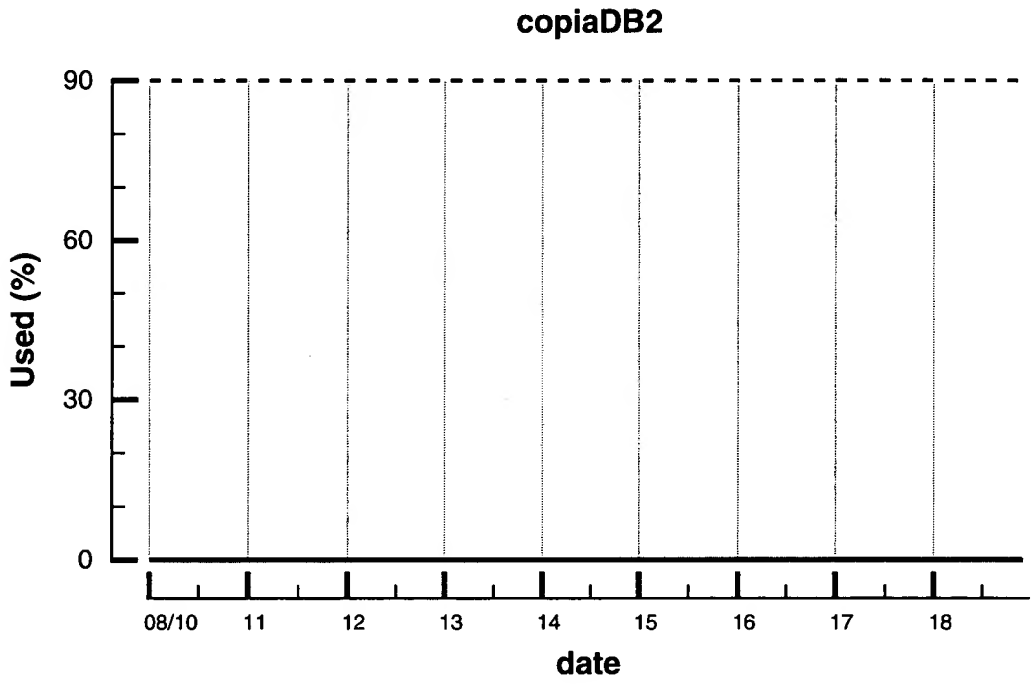

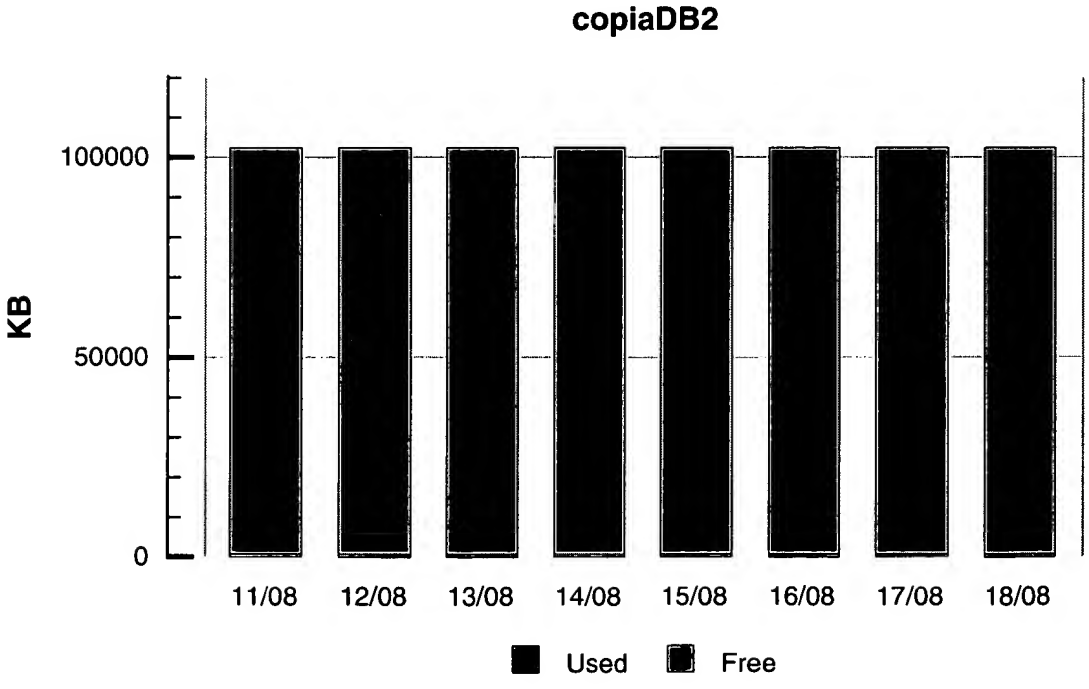

The log usage remained low during the whole monitored period, not indicating any problems.

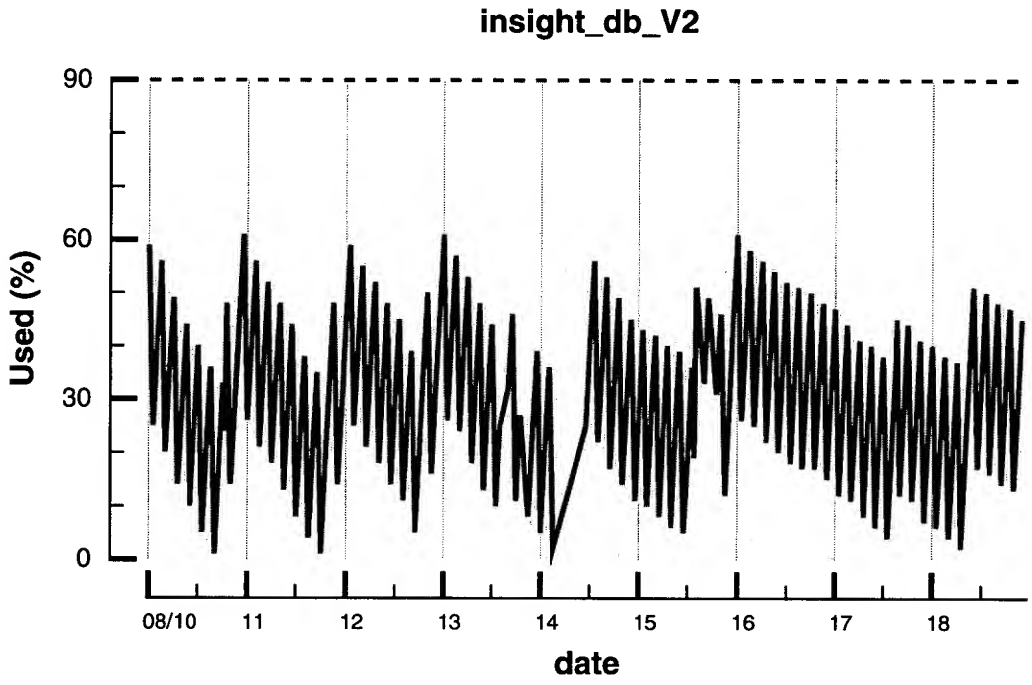

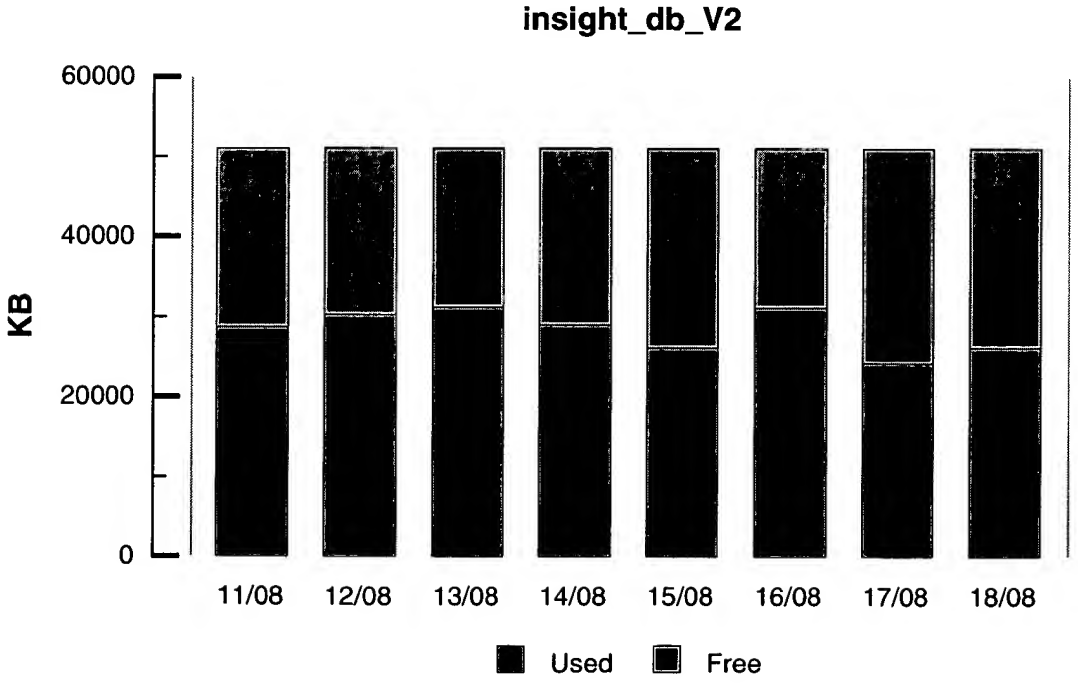

The log usage remained low during the whole monitored period, not indicating any problems.

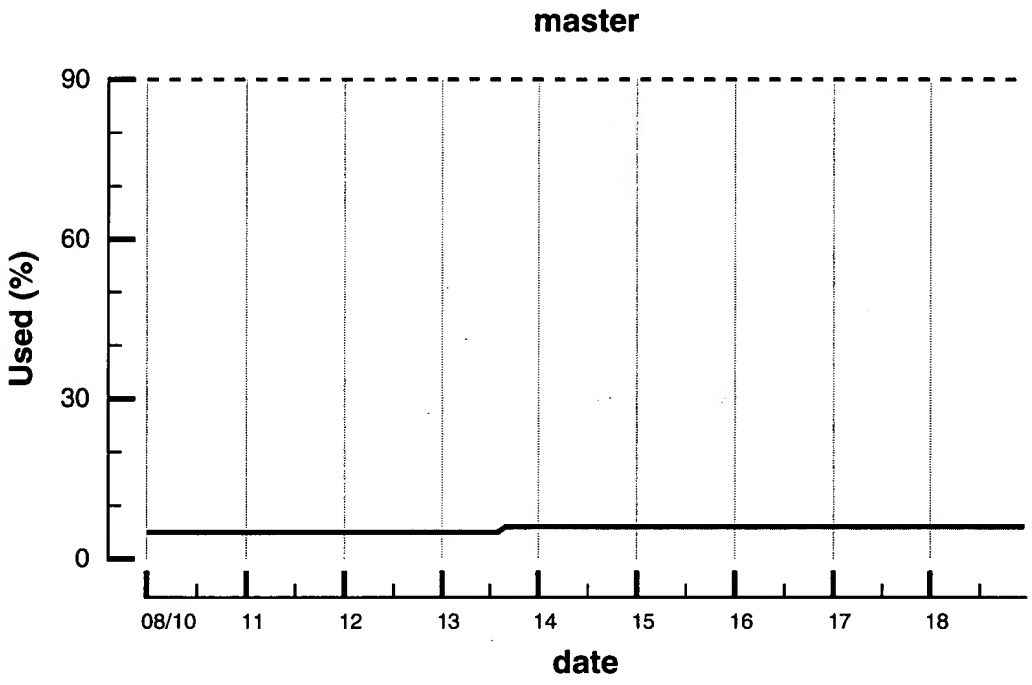

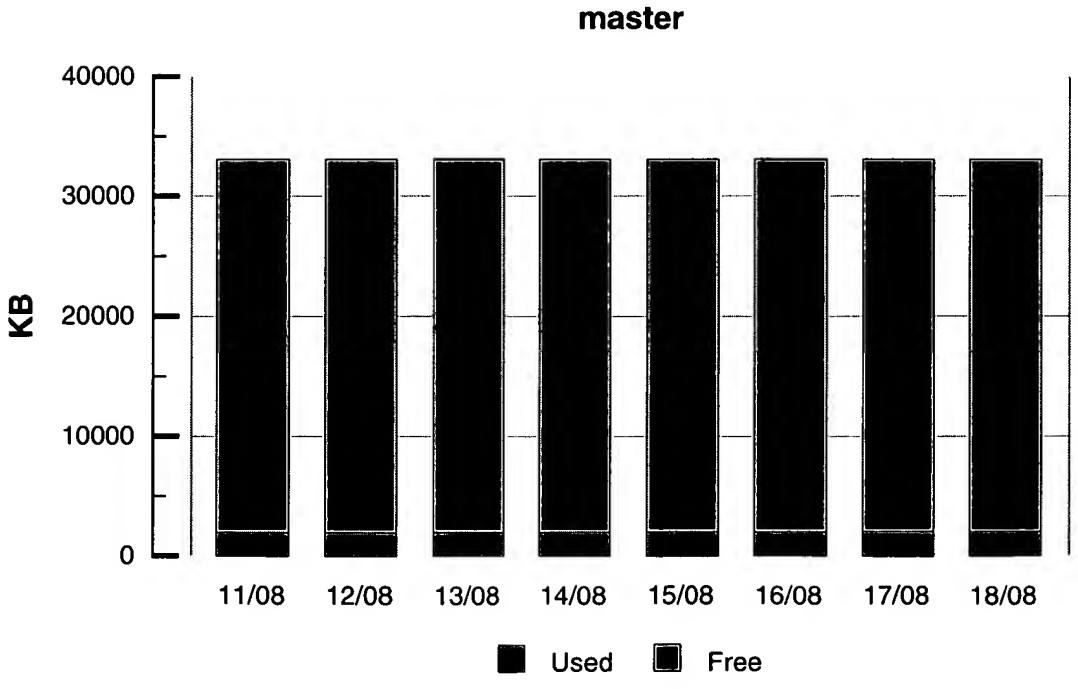

The log usage remained low during the whole monitored period, not indicating any problems.

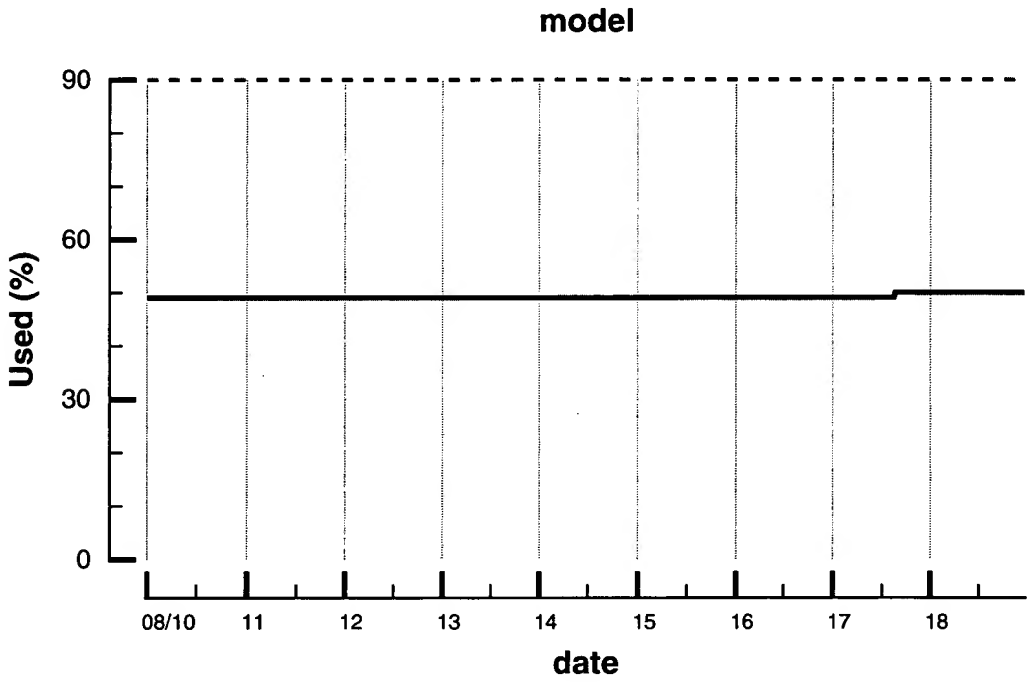

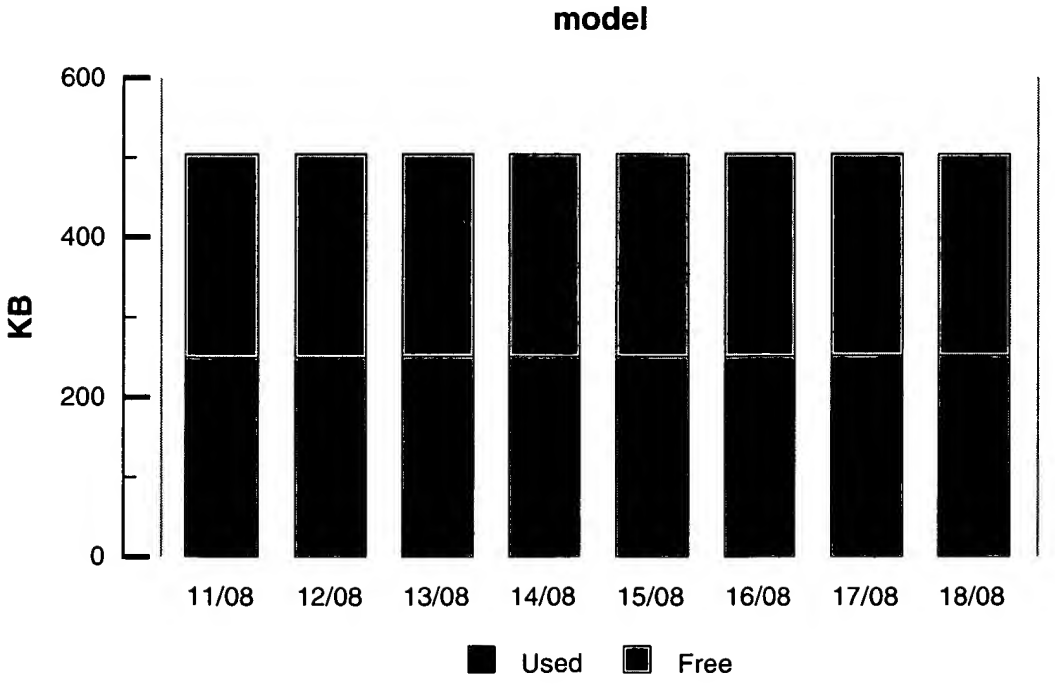

The log usage remained low during the whole monitored period, not indicating any problems.

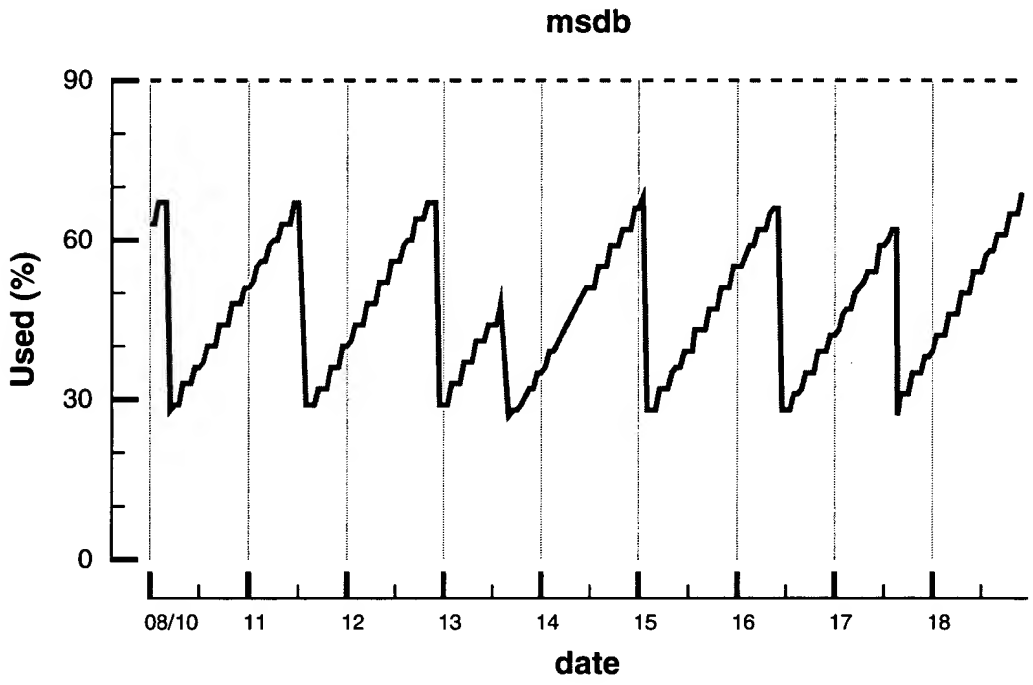

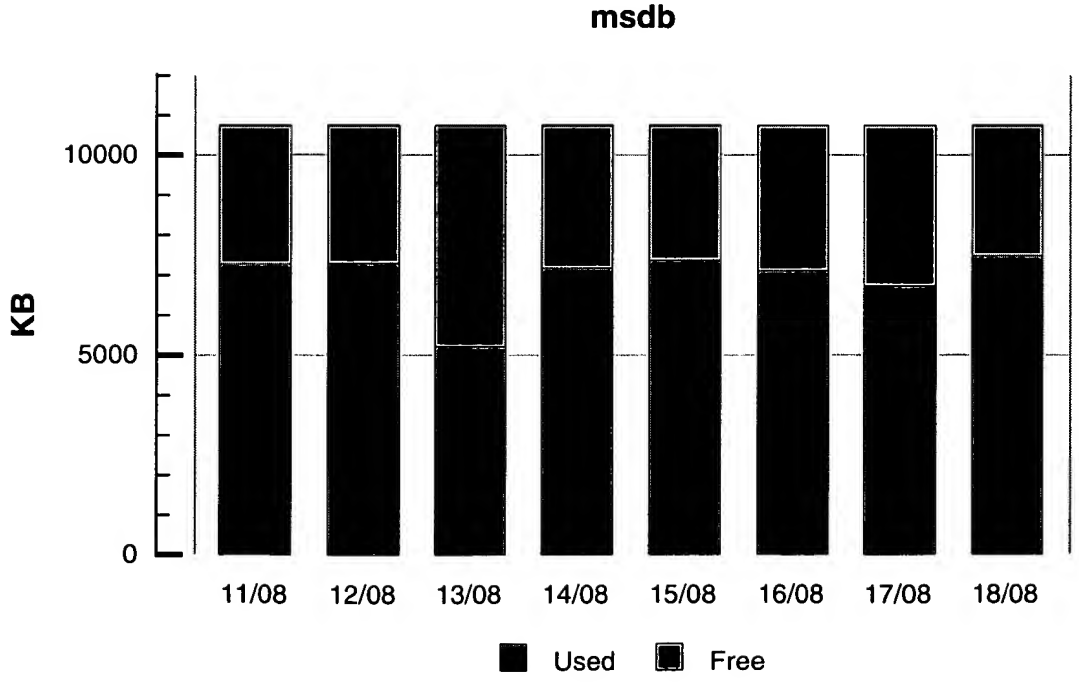

The log usage remained low during the whole monitored period, not indicating any problems.

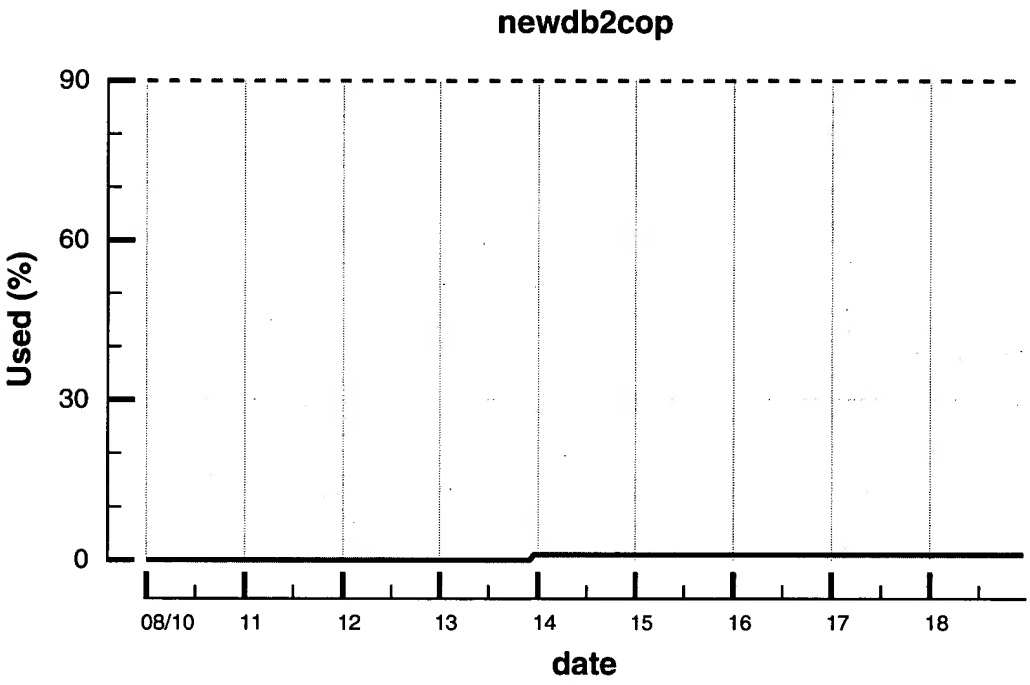

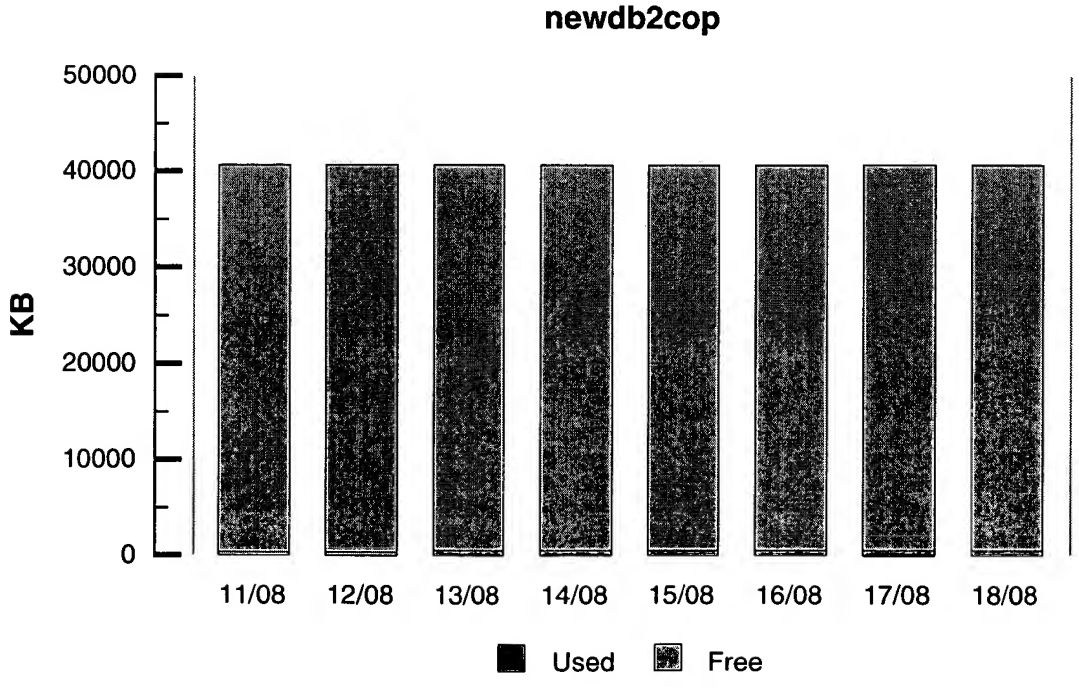

The log usage rate remained high throughout the monitored period, indicating the possibility of a shortage in disk space.

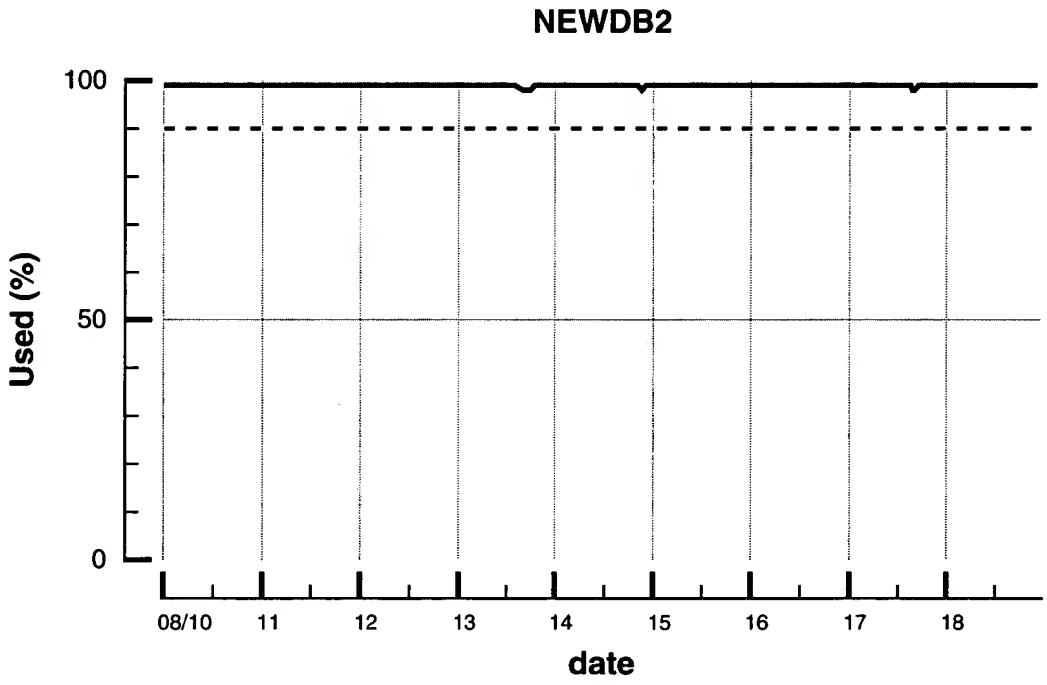

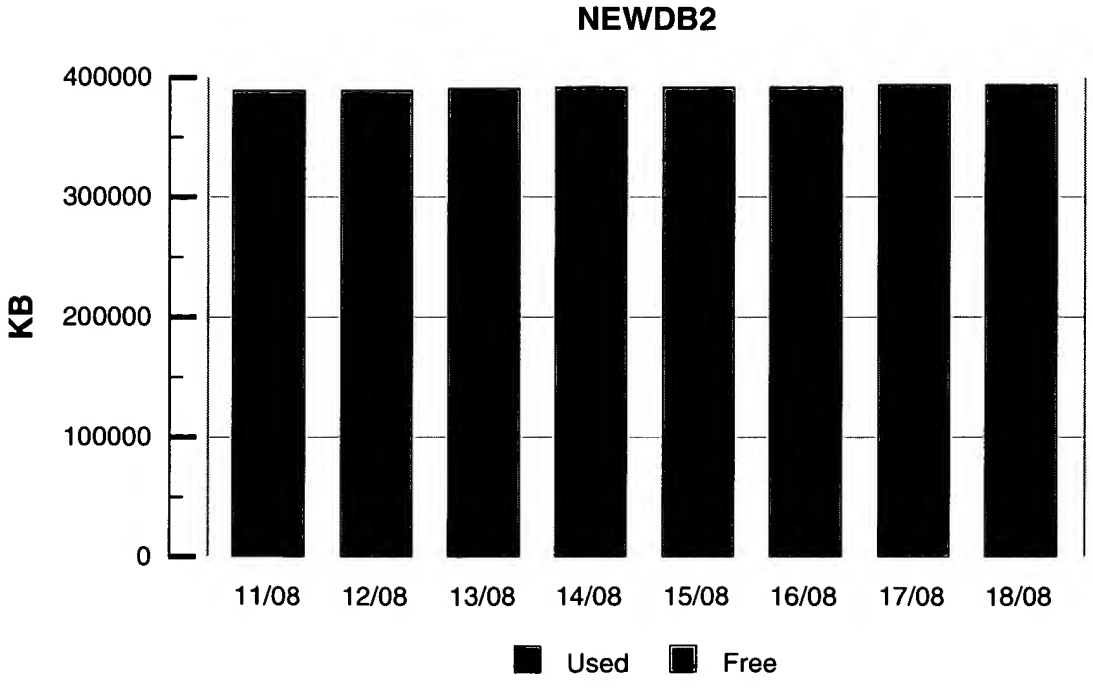

The log usage remained low during the whole monitored period, not indicating any problems.

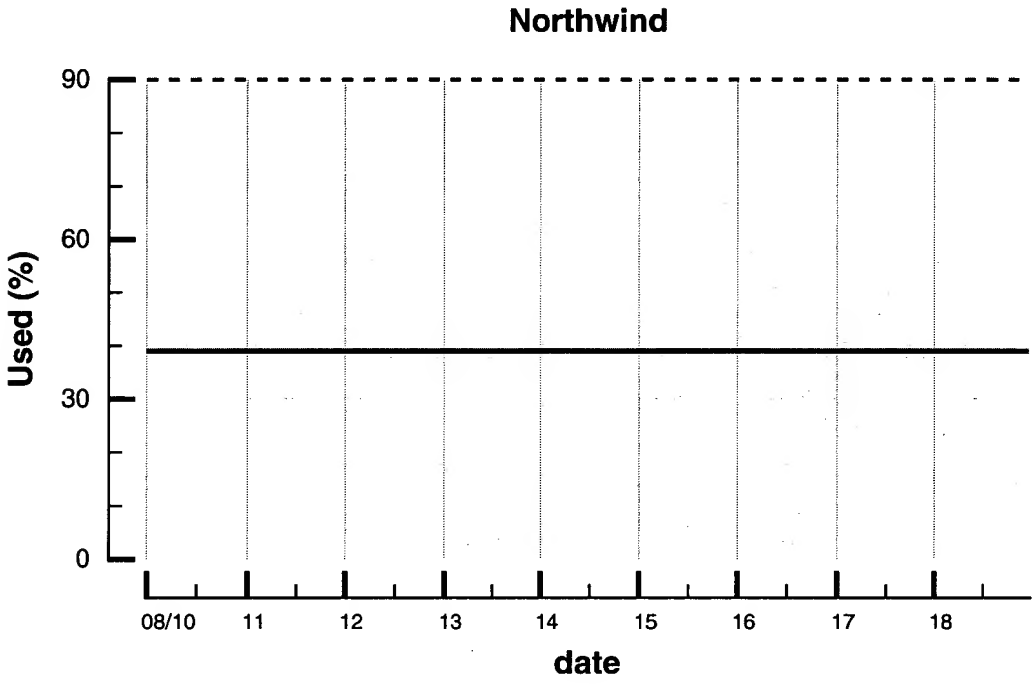

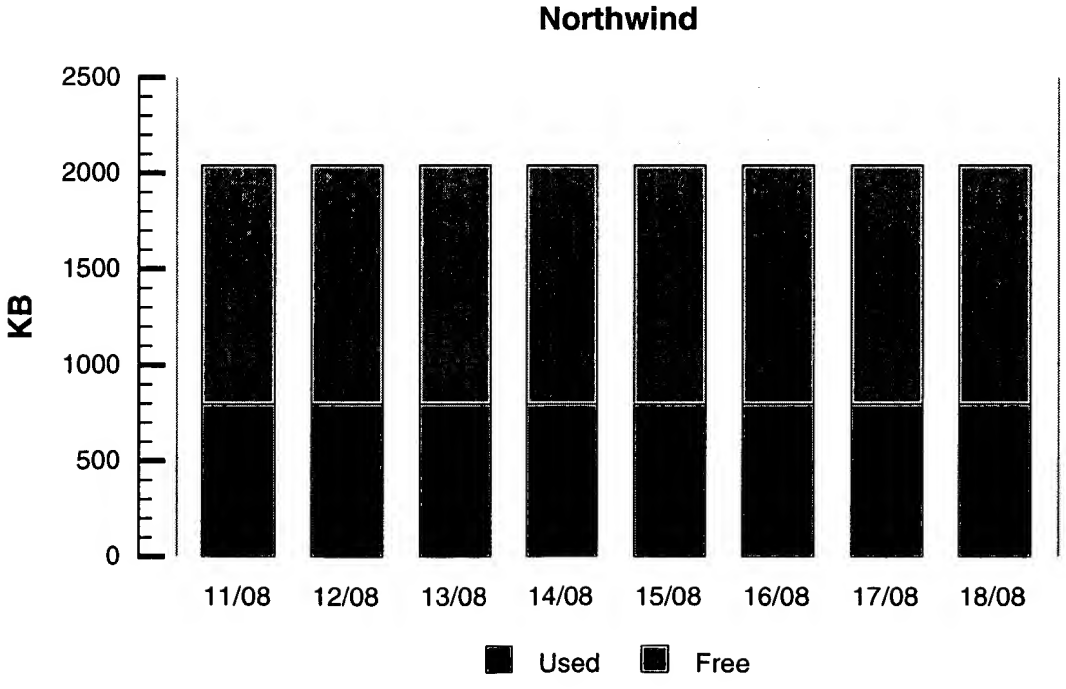

The log usage remained low during the whole monitored period, not indicating any problems.

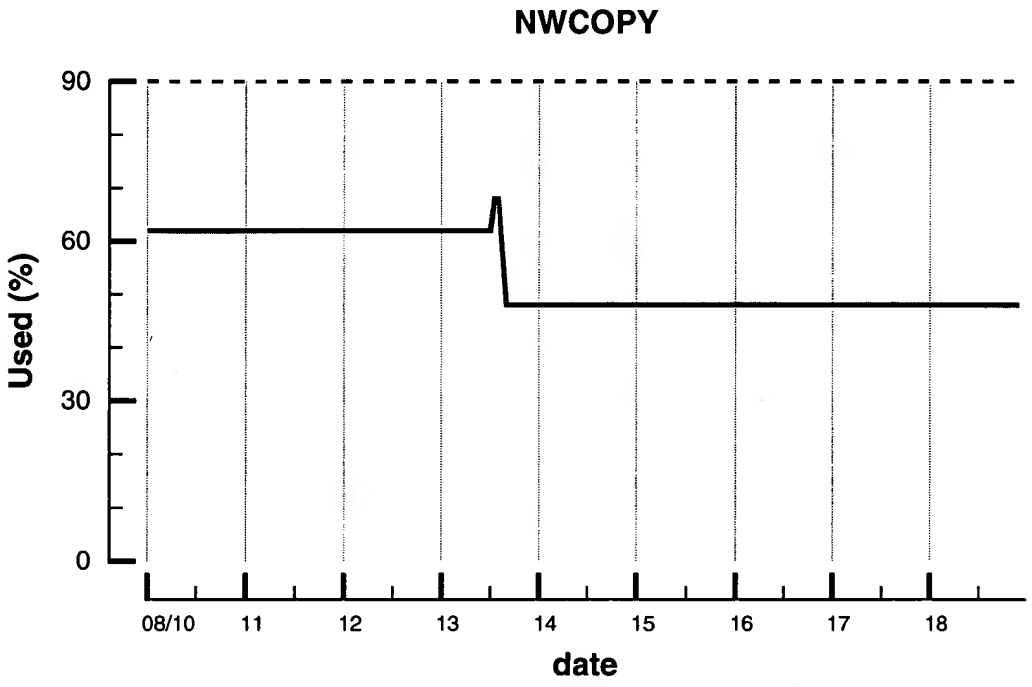

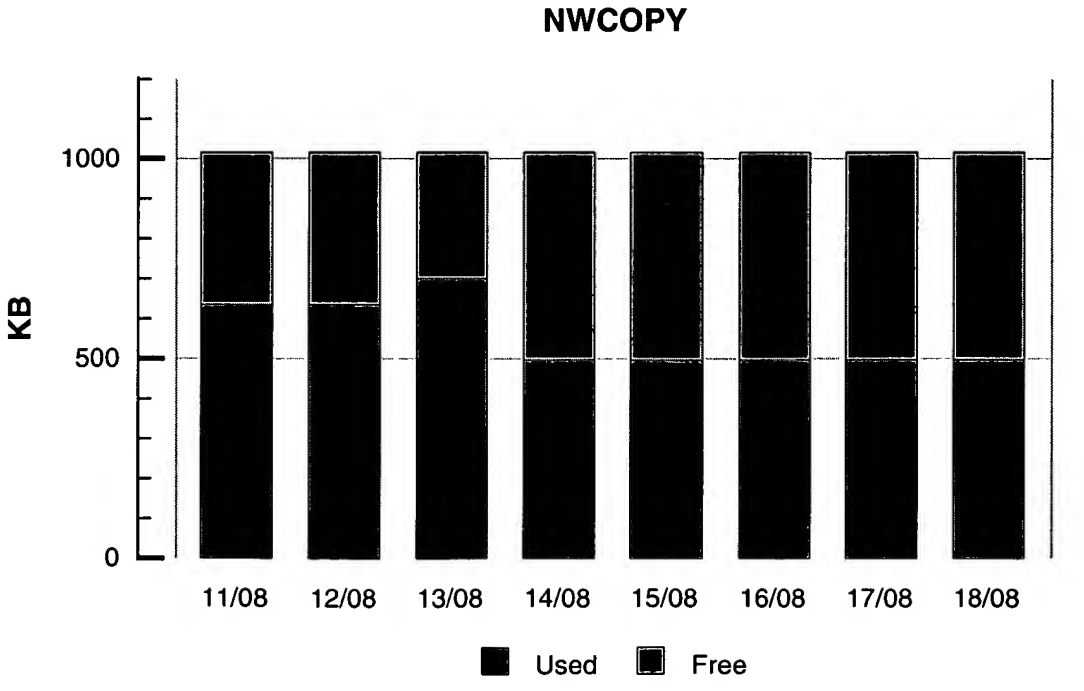

The log usage remained low during the whole monitored period, not indicating any problems.

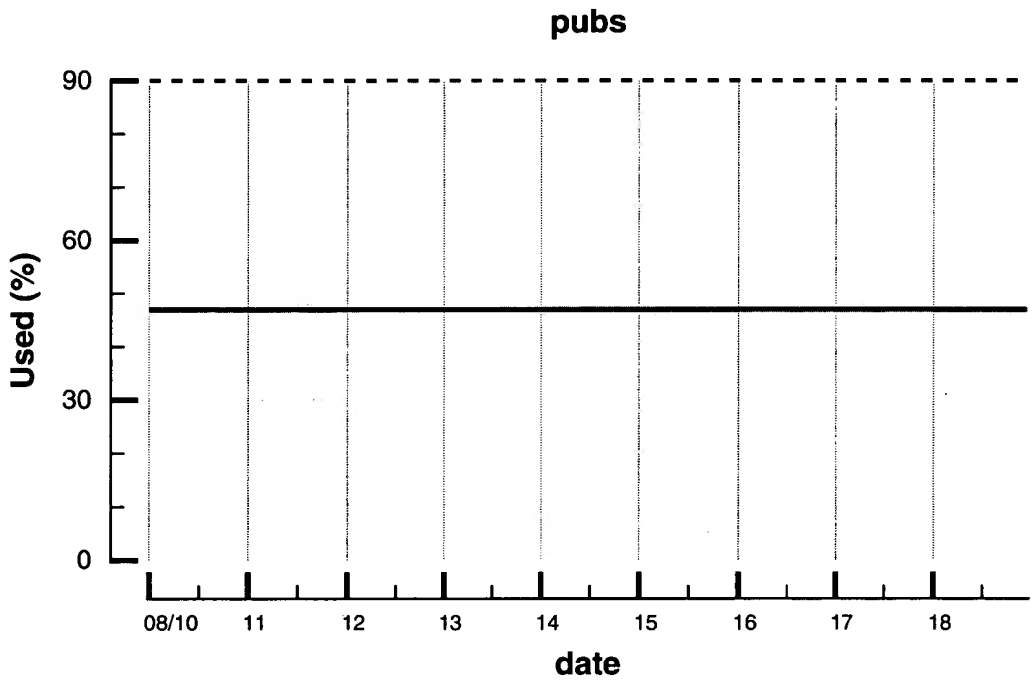

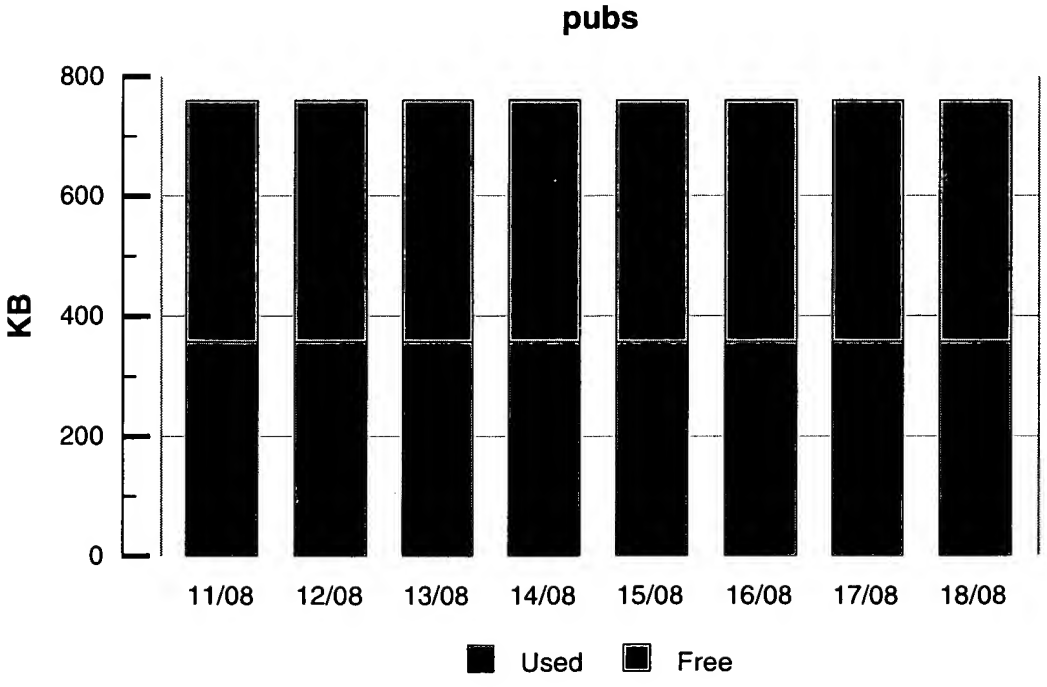

The log usage remained low during the whole monitored period, not indicating any problems.

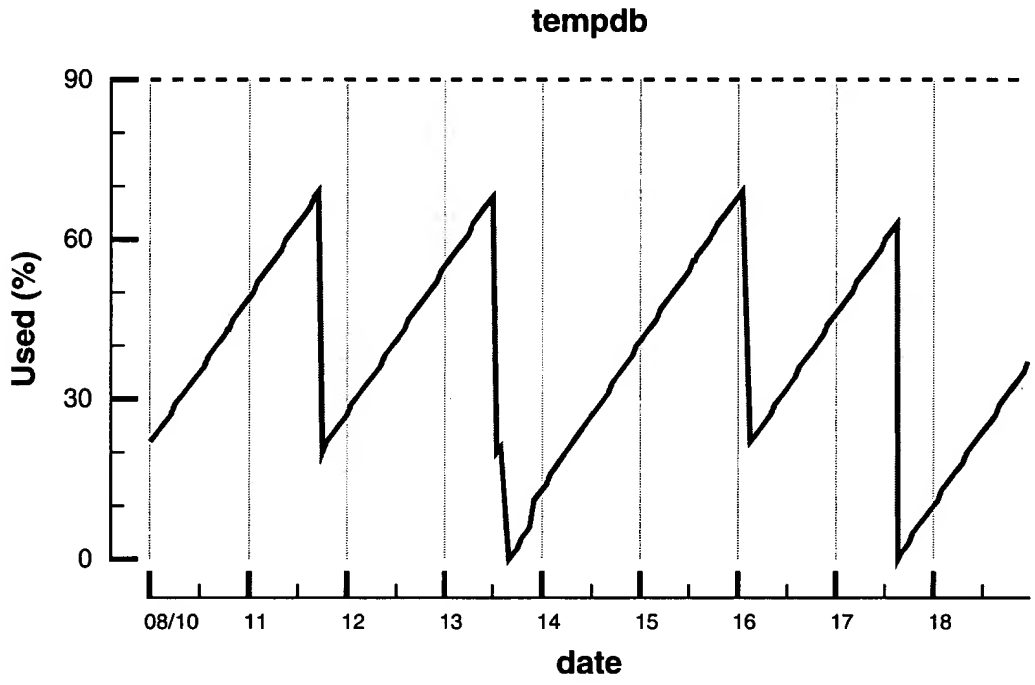

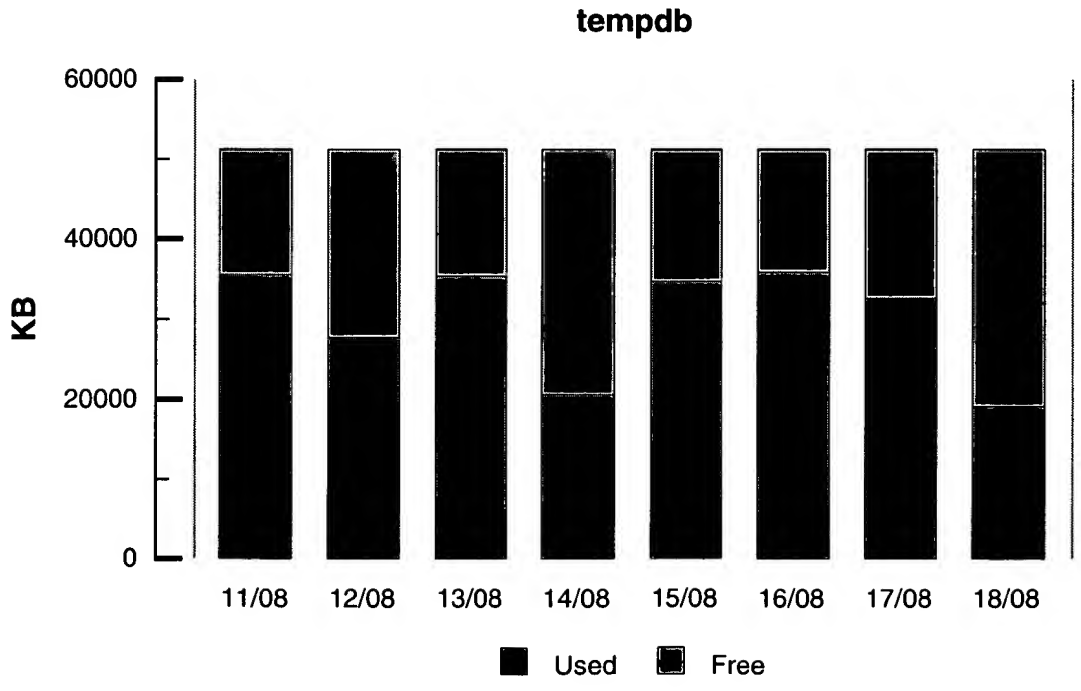

Below is the list of system backups, their name, creation date and ending date.

Name	Creation date	Ending date
copiaDB2	2001-07-30 10:07:35	
insight_db_V2	2001-06-07 20:12:19	2001-07-06 11:38:35
master	2000-08-06 01:29:12	2001-07-22 02:00:34
model	2000-08-06 01:40:52	2001-07-15 02:00:23
msdb	2000-08-06 01:40:56	2001-07-01 02:00:38
newdb2cop	2001-06-27 19:04:26	
NEWDB2	2001-06-25 16:56:44	2001-08-07 18:01:05
Northwind	2000-08-06 01:41:00	
NWCOPY	2001-06-29 12:36:40	2001-07-29 02:00:18
pubs	2000-08-06 01:40:58	
tempdb	2001-08-17 15:19:13	

In order to understand a performance analysis report, one must review a few basic concepts. Performance measurements are arbitrary and are usually based on the perception of the end users and on the nature of the commercial application. In a typical application, end user queries in T-SQL (see Glossary) format are sent through the network to the SQL Server 2000. The SQL Server 2000 executes these queries and, if necessary, sends the results back to the client. The time from the start to the end of this procedure is called the response time of the query. Response time measurement defines the system's performance. If the response time is within a reasonable limit, performance is considered satisfactory. There are several resources and variables that may have a significant effect upon the system's performance. These resources are:

· Memory

· Locks

· Disk occupation

· Number of connections

· Hit rate

· Logs

· CPU

· Disk I/O

Each one of these values, collectively or individually, may affect the performance of the system. Capacity planning, therefore, is an important step in the project of an application. The resources must be carefully evaluated, to make the environment as efficient as possible. Below is a description of the resources and clues to optimize them:

1 - Memory

Memory has different functions in a database environment. It is a fast storage area for program data and disk data caching. Memory consumption, considered not only as real memory (RAM), but also including the virtual memory subsystem, may be evaluated in paging activity, virtual memory usage and paging space. There are a few parameters that indicate the SQL Server 2000's performance, regarding memory usage. From these we may infer if the database occupies too much memory and if there is the possibility of a memory shortage. Target Memory is the total memory reserved by the SQL Server 2000, while SQL Memory indicates how much was really used. These values should be close together, unless the database is idle. Fill factor indicates the percentage of free space on each index page; too large a percentage means the index is fragmented, too small a percentage implies

performance degradation. Normally, the ideal value for the Fill factor is about 50%, but if there are many page splits the Fill factor should be reduced to 30%.

2 - Locks

Locks occur when memory resources are blocked by the SQL-Server 2000 to prevent two users from accessing and altering the same table at the same time. If there are too many locks, the database's performance will deteriorate, since users will constantly have to wait for resources to become available.

3 - Disk Occupation

Indicates if the log file occupies a space proportional to the data file, and how much space it really uses. If the data file does not have the autoextent attribute, it will be necessary to reduce or increase the log file accordingly.

4 - Connection Counter

This expresses the number of connections to the SQL Server in a given period. Too many connections may cause a memory shortage, since each user consumes 40 Kbytes.

5 - Hit Rate

SQL Server 2000 has a buffer cache to improve response time for data that is frequently accessed. If the hit rate of the buffer cache is high, the disks were less used. This improves the performance of the database

6 - Logs

Log monitoring indicates how many alterations are being made in the database. For a database with a lot of recording activity, the time for an alteration to be executed may be reduced to increase security.

7 - CPU

Measuring CPU usage in the SQL Server 2000 is crucial for detecting performance problems. If total CPU usage exceeds 80% for long periods, there is a CPU bottleneck. Also, if the process queue is greater than 1 process per processor, a CPU bottleneck will be identified.

8 - Disk I/O

If the I/O subsystem is working efficiently, every time a server has to read or write data it will do so without waiting. But if a server has too great a workload, reading and writing will have to alternate. This can significantly reduce the performance of the server.

General Configuration Tips for the SQL Server

In this document we will present a description of several important factors in the configuration and monitoring of the database server.

- Partitions formatted in NTFS

This kind of format should only be used up to 80% capacity. Beyond this point, the I/O performance will fall drastically. NTFS partitions need space for management Limit the number of network protocols in Windows If a great number of protocols are configured in the server, an increase of traffic in the network may harm performance. It is recommendable to always use TCP/IP for communication between SQL Server and the clients.

- Sort Order

The sorting order chosen during installation of the SQL Server may affect performance. These are the possibilities:

Binary - The fastest, but may cause problems in the client applications.

Case Insensitive - The second fastest, use it if possible.

Accent insensitive, uppercase preference and Case Sensitive - The slowest

- Location of the Data Files

The best procedure is to create the data files and logs in separate disks or arrays, with the finality of isolating reading and writing conflicts. Use the most available physical disks for creation of the data files.

- Max assinc I/O

If the SQL Server has an excellent disk controller, this parameter may be increased. The default value is 32 (maximum is 255), which is sufficient in most cases. The initial rule when increasing this parameter is multiplying by 2 or 3 the number of physical drives that support simultaneous I/O.

- Recovery Interval

This parameter defines the appropriate recovery interval for the SQL Server. If a server is actively being used for INSERT, UPDATE and DELETE operations, it is possible that the default value for this parameter (0) is not ideal. If the server presents periods with 100% read/write activity, this value may be increased until an optimum value

is found.

- SP_TABLEOPTION PINTABLE

If a small table (or some) is used much more than the others, this parameter may be used to keep it always in the cache after its first reading.

- Network Packet Size

If the information switched between the clients and the servers are images, or any other large piece of data, this parameter may be increased to improve performance (default is 4096 bytes.

- Max Degree of Parallelism

If the SQL Server is used in OLTP and not OLAP applications, this parameter may be disenabled to increase performance of the server. When this parameter is enabled, the SQL Server analyzes every query to verify the possibility of dividing it in more than one processor. This is unnecessary in OLTP applications, since in this case most operations are simple and do not require parallelism.

- Max Worker Threads

This parameter indicates the maximum number of threads that the server reserves for the SQL Server (sqlservr.exe). Each user connection uses one thread. If there are more connected users than available threads, the SQL Server will use thread pooling, degrading performance. This parameter (default 255) will always be slightly greater than parameter Max user connections.

- TempDB

If the database is being used a lot by the applications, its physical location should be apart from other data files.

- Comments about Data Files and File groups

1. A data file or a file group may not be used in more than one database.

2. A data file may only belong to one file group

3. Data and log may not belong in the same data file.

4. File groups are mechanisms used to associate objects to specific files.

5. Tables may only belong to one file group.

6. Several data files may be created in different disks and associated to a

single file group.

7. File groups have a proportional growth strategy- the free space in the data files will always be proportional. For example: if a data file has 100 MB and another 200 MB, for each byte recorded in the first data file, two will be recorded in the second.

- Monitoring of the Data Files

Normally, the more data files, the better parallelism will be, but in some data files the saturation point may have been reached. To evaluate how often the data files are used, we must use the Windows NT Performance Monitor, checking the Physical Disk e Disk Queue Length counters.

Extended Memory Size

This option is available in the SQL Server for future versions of the Windows NT on alpha platforms. It indicates the number of megabytes that the SQL Server will use as cache in the memory, above 2GB.

- Glossary

Buffer cache - A definite amount of physical memory that is reserved by SQL -Server for data that is used more frequently. This reduces disk usage and improves general performance.

Disk I/O - This process strongly affects performance, since read/write operations are usually much slower than physical memory operations.

Index page - Index page used by SQL-Server to ease the search for information

Log file - File where all database alterations are registered. It is possible to restore information if there are problems.

Paging - Memory is organized in pages. The operating system can transfer these pages from physical memory to disk and vice-versa. This process is called paging.

Process queue - This is the queue of processes "waiting" to be processed by the CPU.

T-SQL - Transact-SQL is SQL-Server's programming language.

Virtual memory - The operating system manages the total available memory, composed of RAM and disk, creating a single "virtual" memory block.

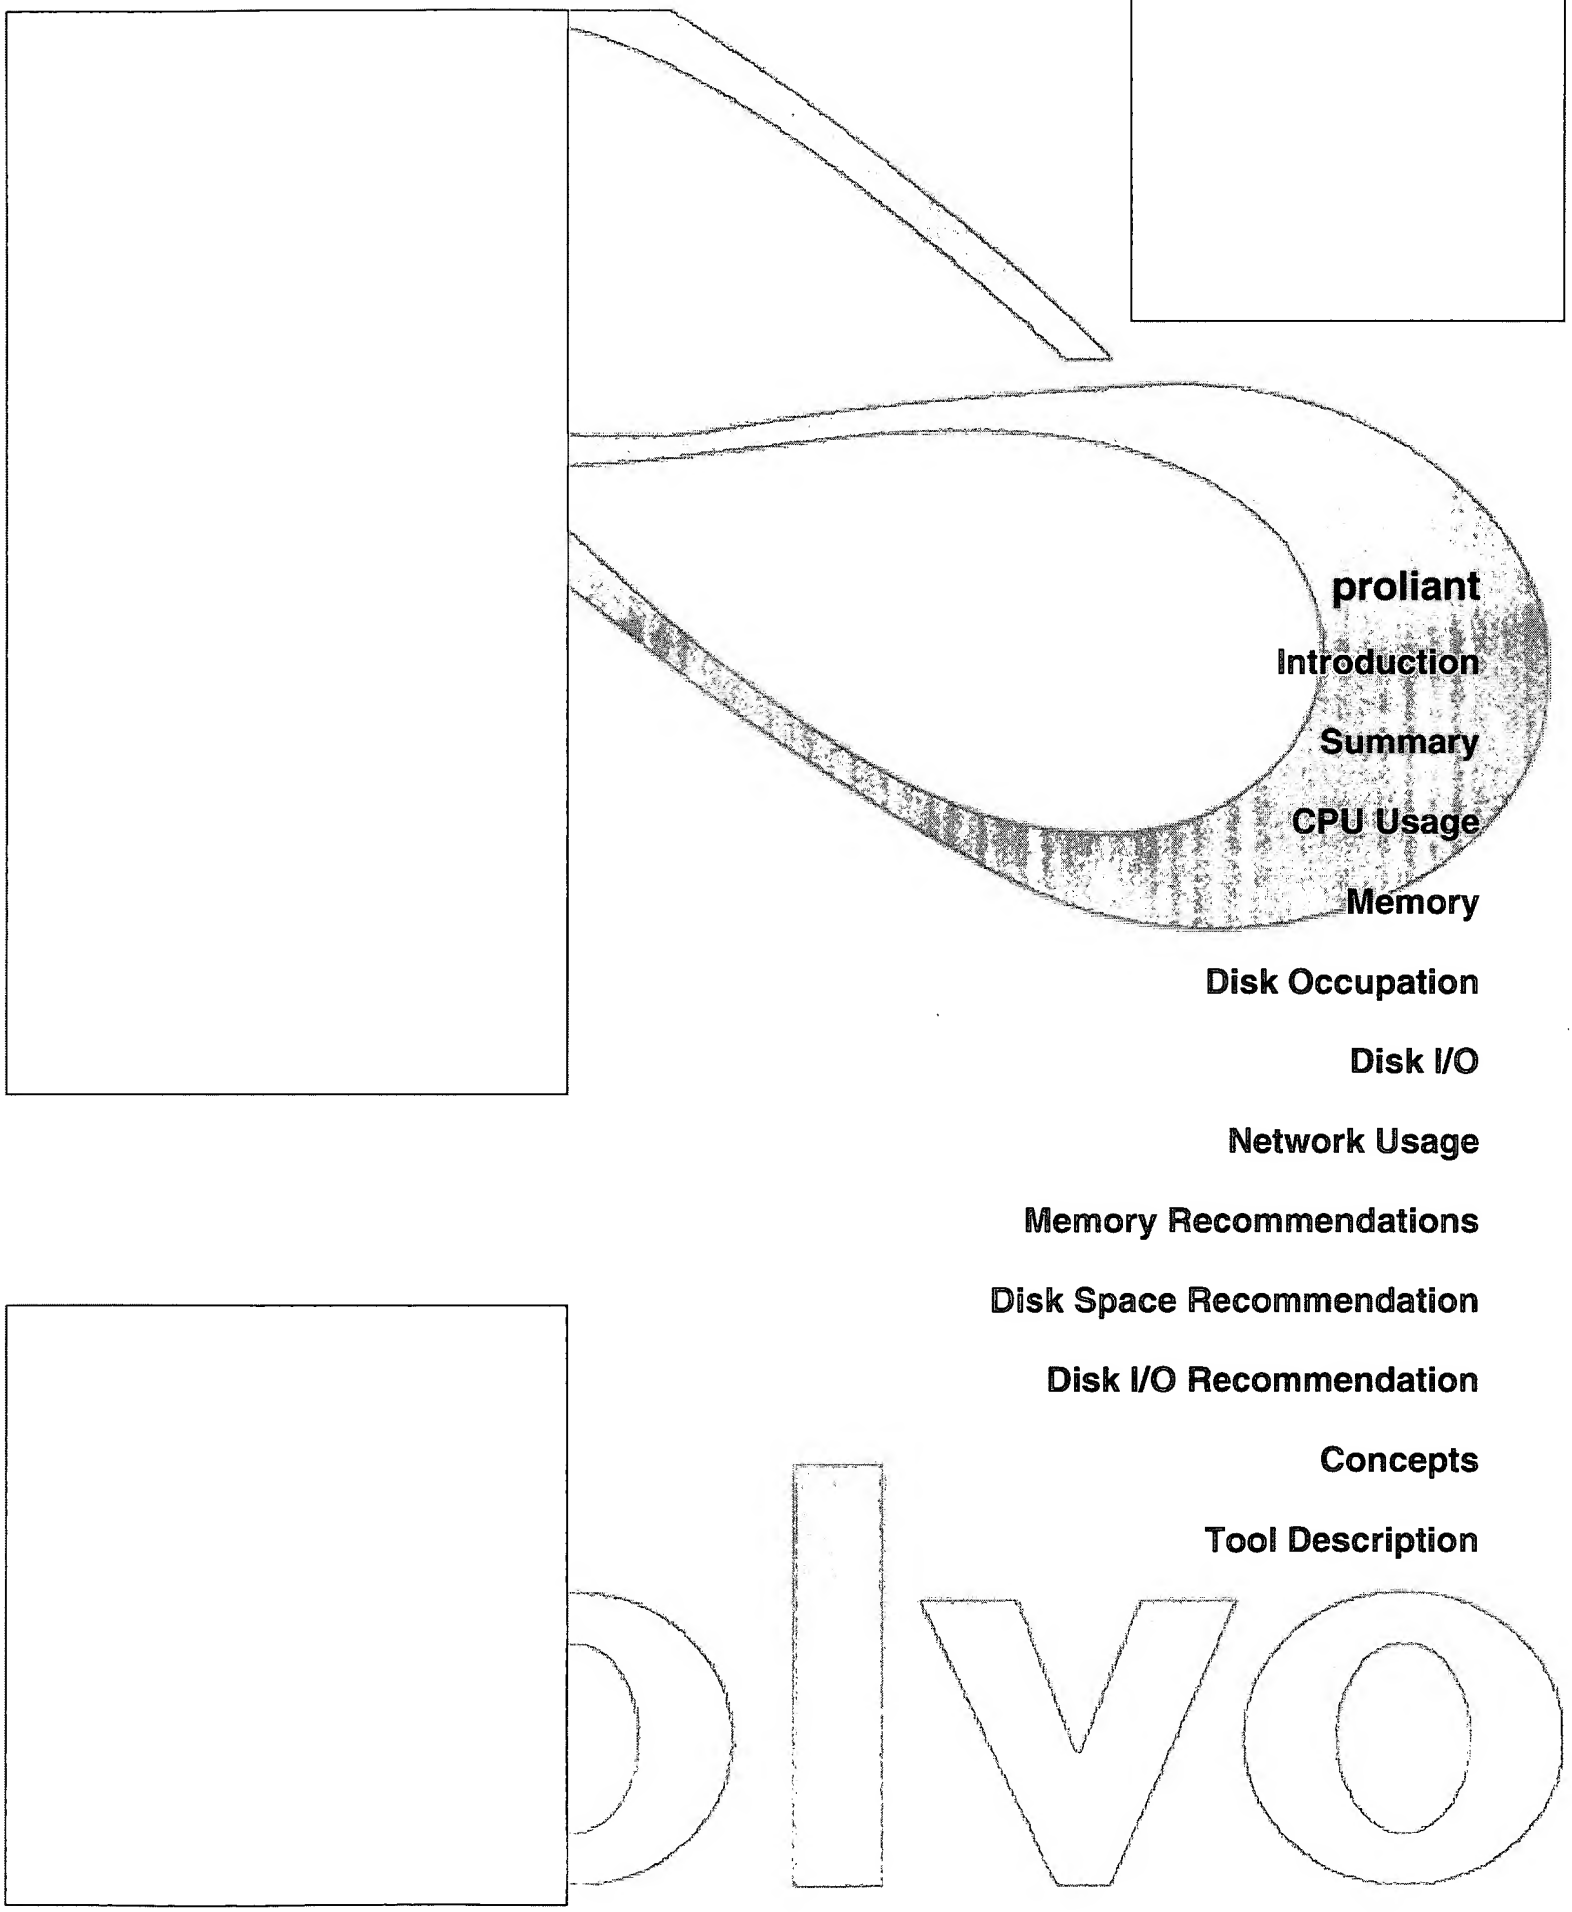
proliant
Introduction
Summary
CPU Usage
Memory
Disk Occupation
Disk I/O
Network Usage
Memory Recommendations
Disk Space Recommendation
Disk I/O Recommendation
Concepts
Tool Description
I/O

Introduction

Based on the data collected in machine proliant, from 08/09/2001, at 00:00, to 09/01/2001, at 17:00, this capacity planning report was produced.

The data used in this report was obtained from an exclusive collector, with high resolution and low intrusion, developed especially for this end. This collector obtained data directly from the core of the operating system, with no need for libraries or additional utilities, and minimum impact on the environment. The collected data is stored in binary format, to provide persistence. The data is automatically sent, compressed and encrypted to ensure fast transit and confidentiality.

The content of this report is the result of years of experience in performance analysis and capacity planning. The tool used to generate this document operates in a totally automatic manner, without direct human intervention. It uses an extensible inference machine, based on heuristics and rules which are continuously improved. Through the use of concepts such as "watermarks"and regression, it is possible to determine when a computational resource will reach its saturation point.

Profile description:

ML330

We have assumed that the workload is CPU-bound, therefore the upgrades will be calculated based on SPECint.

During the monitored period, this was the summary configuration of the target machine:

OS	:	MS Windows 2000 Advanced
Version	:	5.0.2195 (sp 2.0) Service Pack 2
Host	:	proliant
IP address	:	192.168.1.18
Processors	:	1 Pentium III (Coppermine)
Speed	:	728 MHz
Memory	:	191 MB

The last boot in machine proliant was on 08/30/2001, at 19:03.

This report refers to the monitoring which took place between 08/09/2001, at 00:00, and 09/01/2001, at 17:00. The future horizon considered was 180 days. See below this period's highlights:

The CPU had a usage reduction. That is why no projection was made.

Memory is highly saturated, exceeding the limit of 90%, peaking at 228%, with a growth of 20.5% per month. If the amount of memory is increased in 291%, maximum usage should remain below the threshold of 90%, for the considered horizon.

The analyzed disks presented a good occupation, with the average usage not reaching the saturation level of 70%, growing 5.9% per month. The usage peaked at 60.5%. It is estimated that the saturation level will be reached on 01/2002. If the disk space is increased in 4.2%, maximum usage should remain below the threshold of 70%, for the considered horizon

The analyzed disks presented a good I/O performance (time active), with the average usage not reaching the saturation limit of 40%, growing 69.9% per month. The total usage peaked at 32%. It is estimated that the saturation limit will be reached on 11/2001. If the number of disks is increased in 77.9%, maximum usage should remain below the threshold of 40%, for the considered horizon.

Network usage was satisfactory, with the average usage not reaching the saturation limit of 70%. Usage peaked at 0.2% and grew 102.3% per month.

For this environment to operate satisfactorily in a future horizon of 180 days, it is necessary to add 832 MB memory, add at least 0.4 GB of disk space and add 1 disks, spreading the load over them.

CPU Usage

The CPU had a usage reduction. Because of that, there is no projected growth. The highest point was 6.9% (SPECint_rate 22.11). The reliability of the linear regression is of 98.9%.

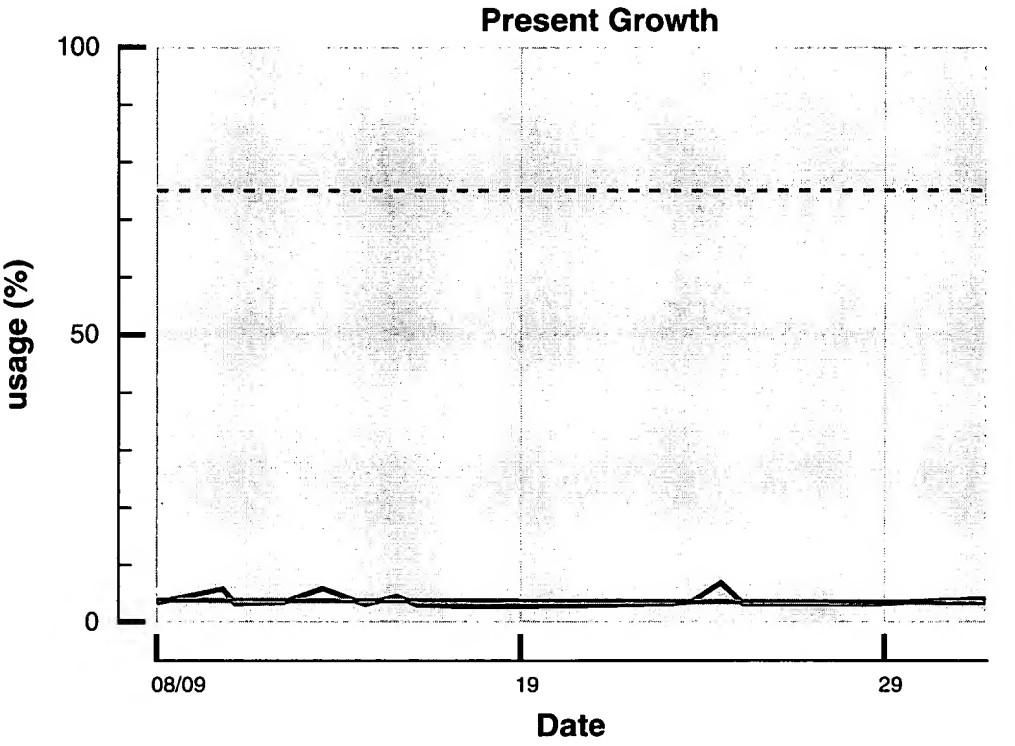

Memory

During the analyzed period, the memory was highly saturated, exceeding the limit of 90%, with a growth of 20.5% per month, peaking at 228%. The future horizon considered is 180 days. If the memory is increased in 291%, the usage will remain below the threshold of 90%, for the considered horizon. The reliability of the linear regression is of 85.9%.

The lines below represent the total memory consumption and its growth.

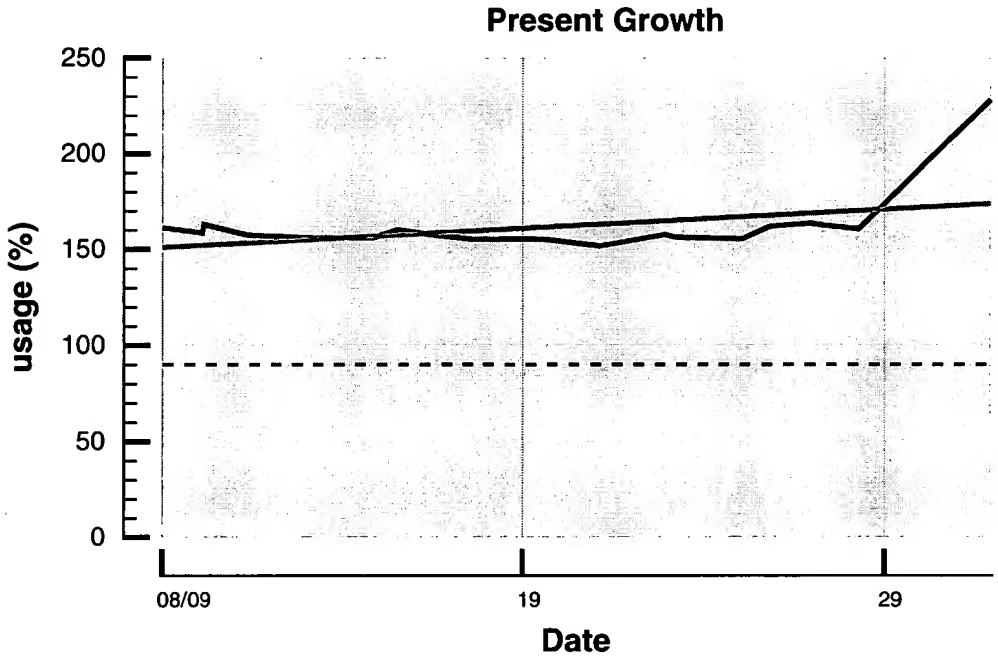

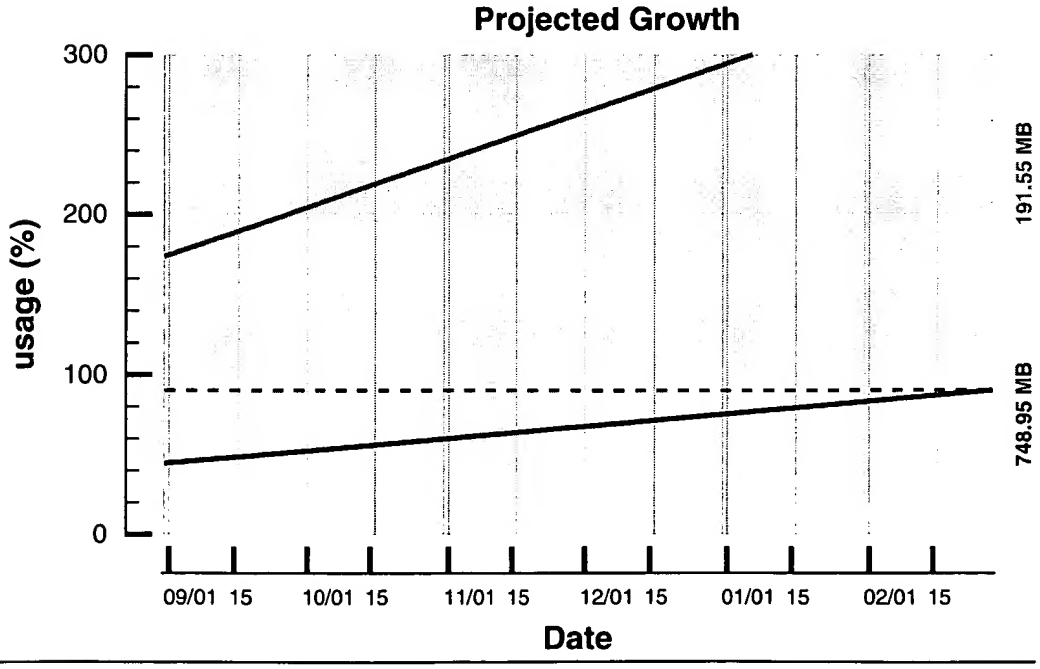

Disk Occupation

During the analyzed period, the analyzed disks presented an average (space) occupation below the saturation threshold of 70%, growing 5.9% per month. The peak was 60.5%. For a future horizon of 180 days, it is estimated that this saturation threshold will be reached on 01/2002. If the disk space is increased in 4.2%, maximum usage should remain below the threshold of 70%, for the considered horizon The reliability of the linear regression is of 98.6%.

All disks are considered here as a single storage device.

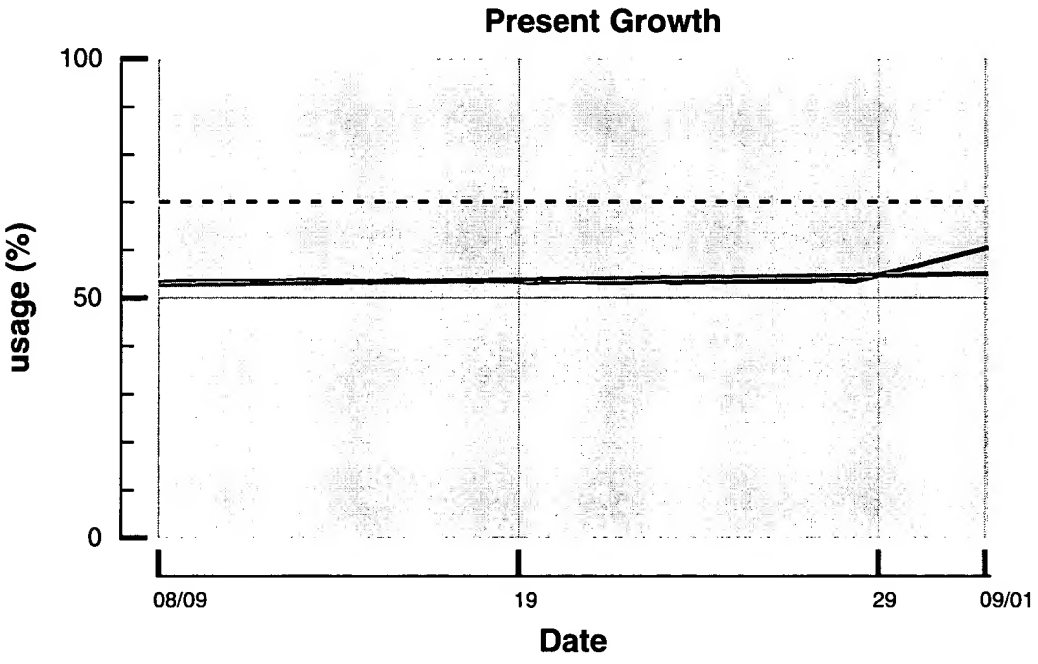

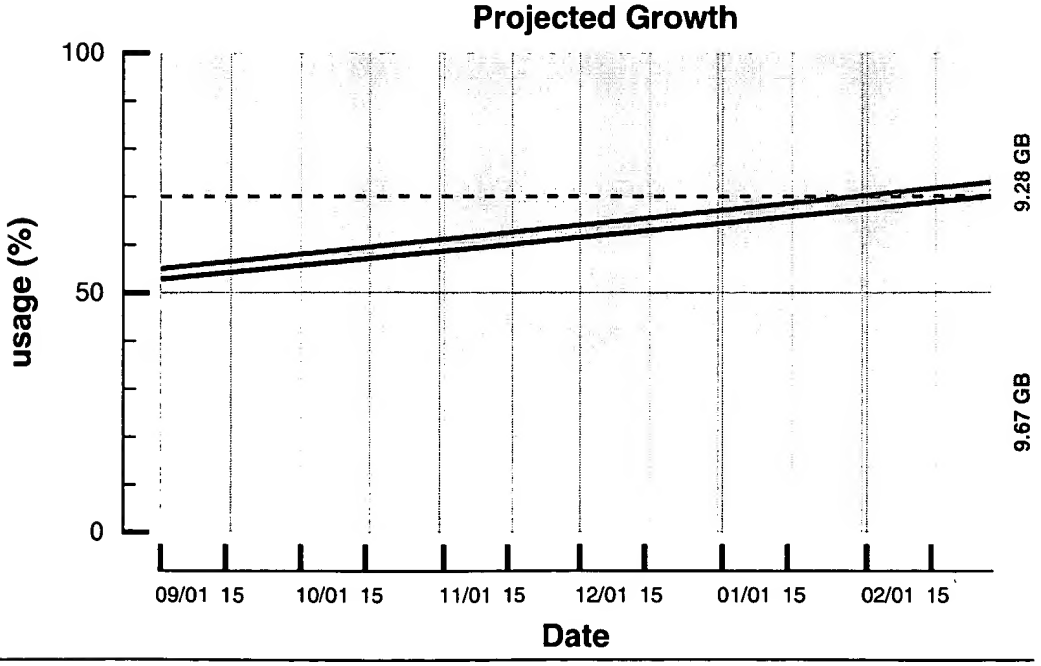

Disk I/O

During the analyzed period, the analyzed disks presented a good I/O performance (time active). The average usage did not reach the saturation limit of 40%, growing 69.9% per month. The usage peaked at 32%. In a future horizon of 180 days, it is estimated that the saturation limit will be reached on 11/2001. If the number of disks is increased in 77.9%, maximum usage should remain below the threshold of 40%, for the considered horizon. The reliability of the linear regression is of 96.4%.

The line below represents the total disk time active. Here, all disks are considered as one.

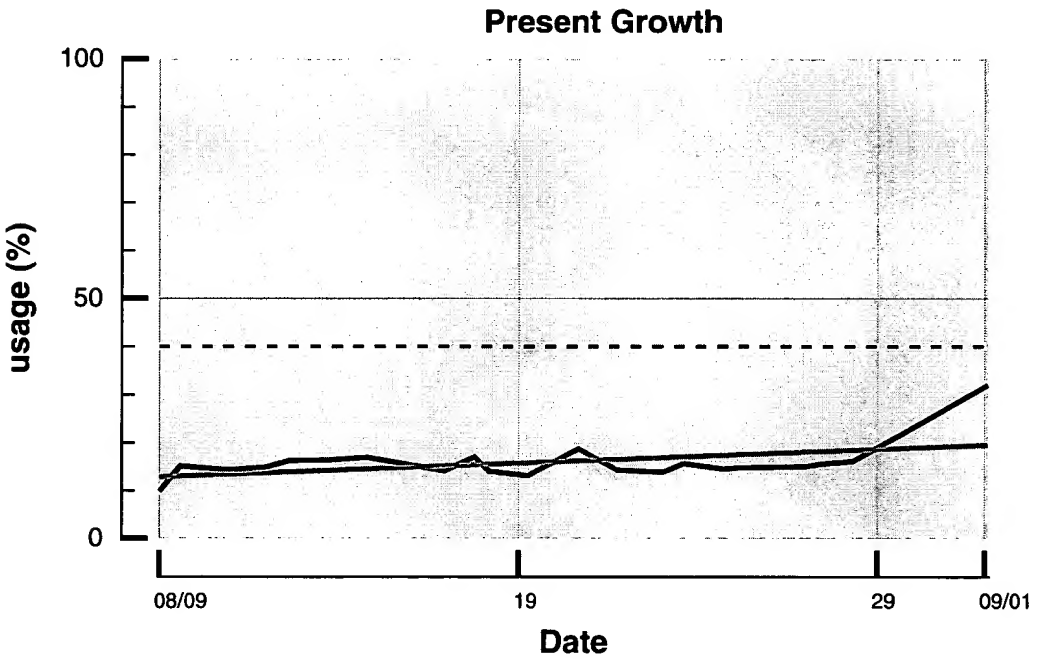

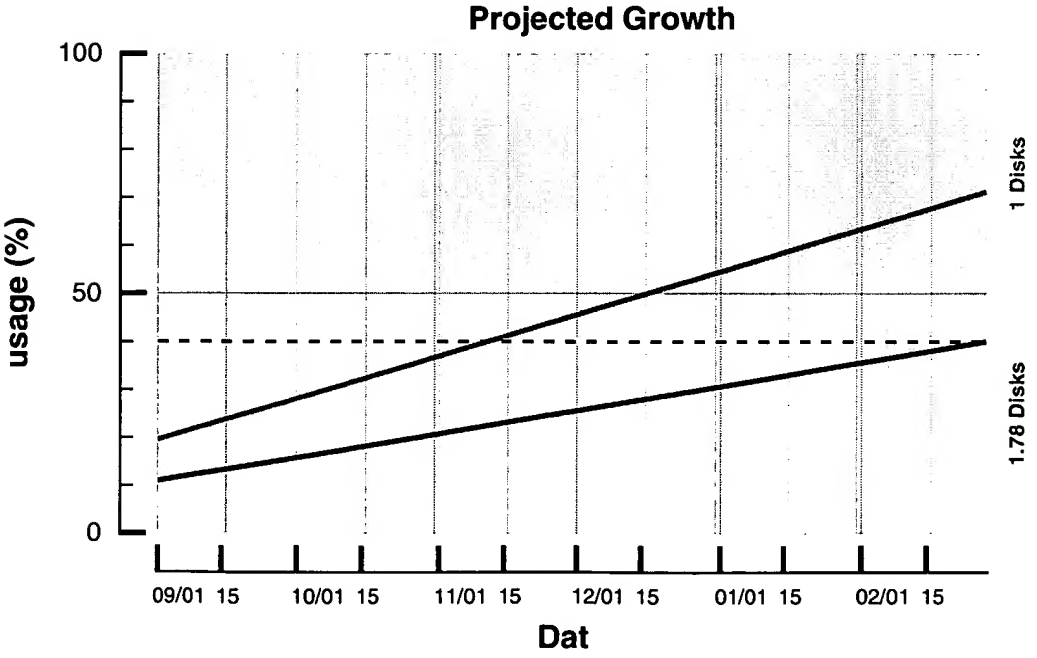

Network Usage

Network usage was satisfactory, with the average usage not reaching the saturation limit of 70%, growing 102.3% per month and peaking at 0.2%. The future horizon considered is 180 days. The reliability of the linear regression is of 100%.

Here the total network bandwidth is considered, and the total consumption, aggregating all network adaptors.

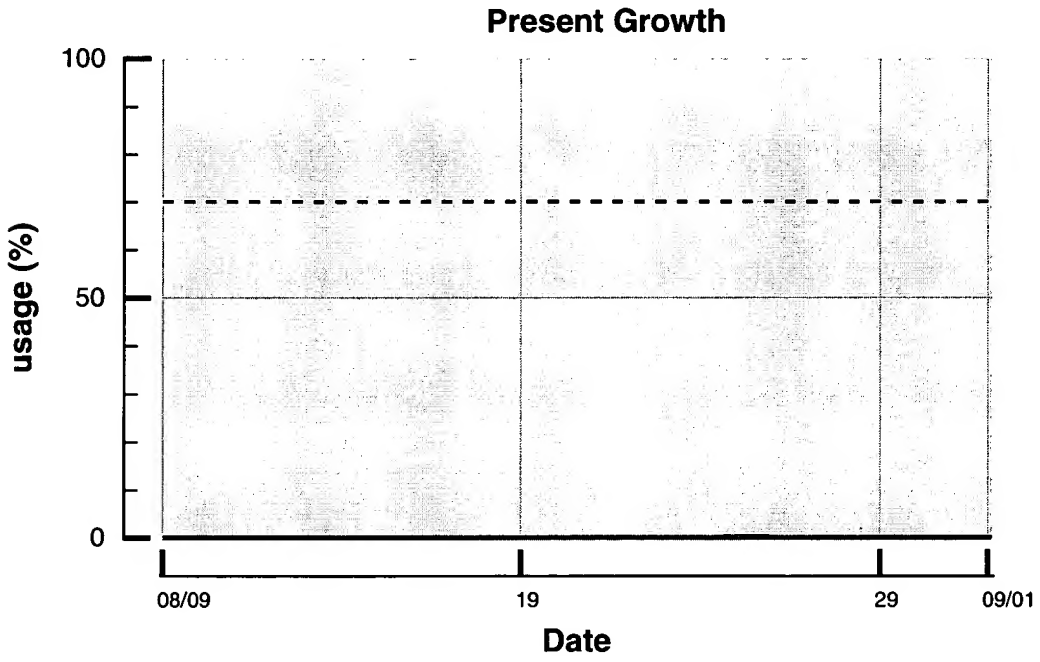

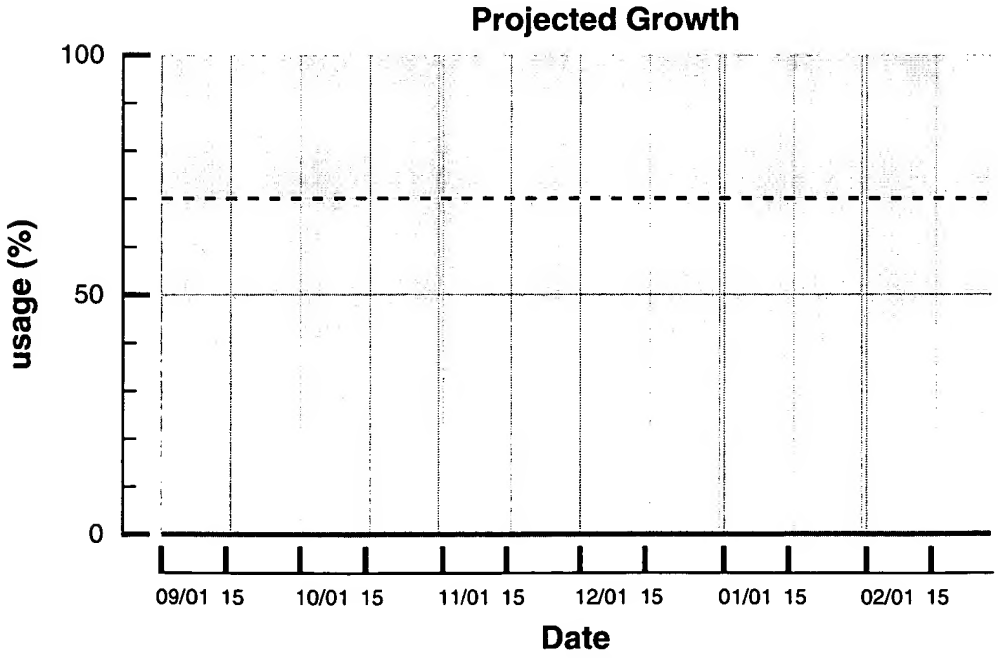

For this environment to operate satisfactorily in a future horizon of 180 days, it is necessary to add 832 MB memory, as shown in the graph below.

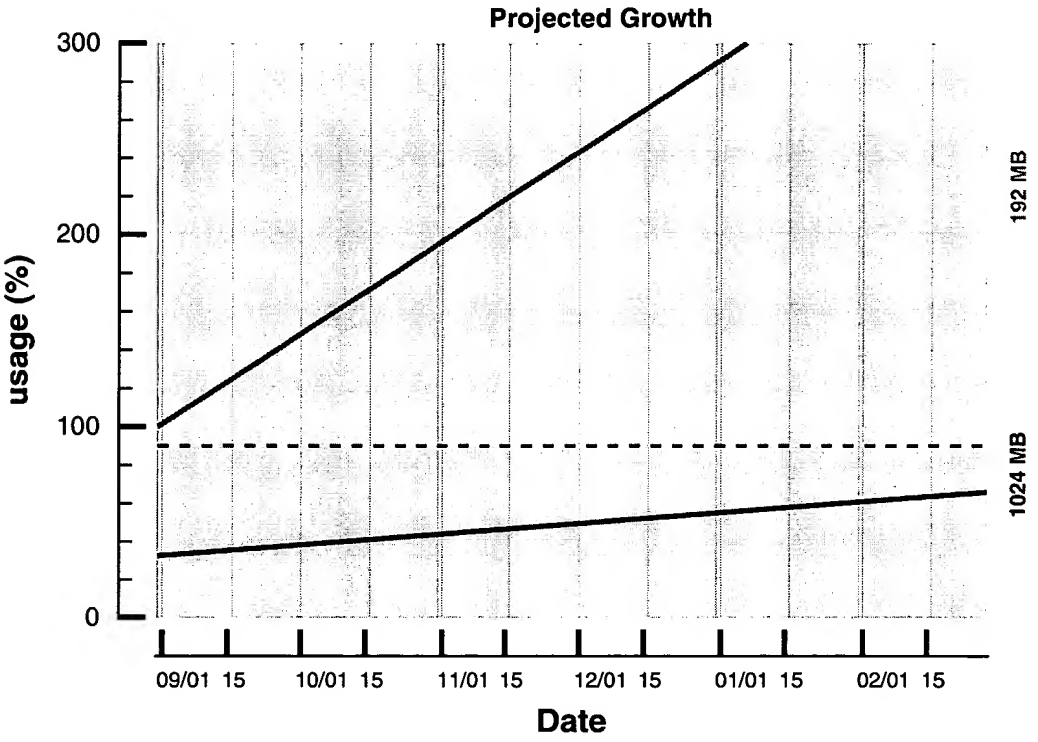

Disk Space Recommendation

For this environment to operate satisfactorily in a future horizon of 180 days, it is necessary to add at least 0.4 GB of disk space, as shown below.

The lines represent, respectively, the present configuration (orange) and the recommended configuration (brown) for the machine to operate within the usage limit.

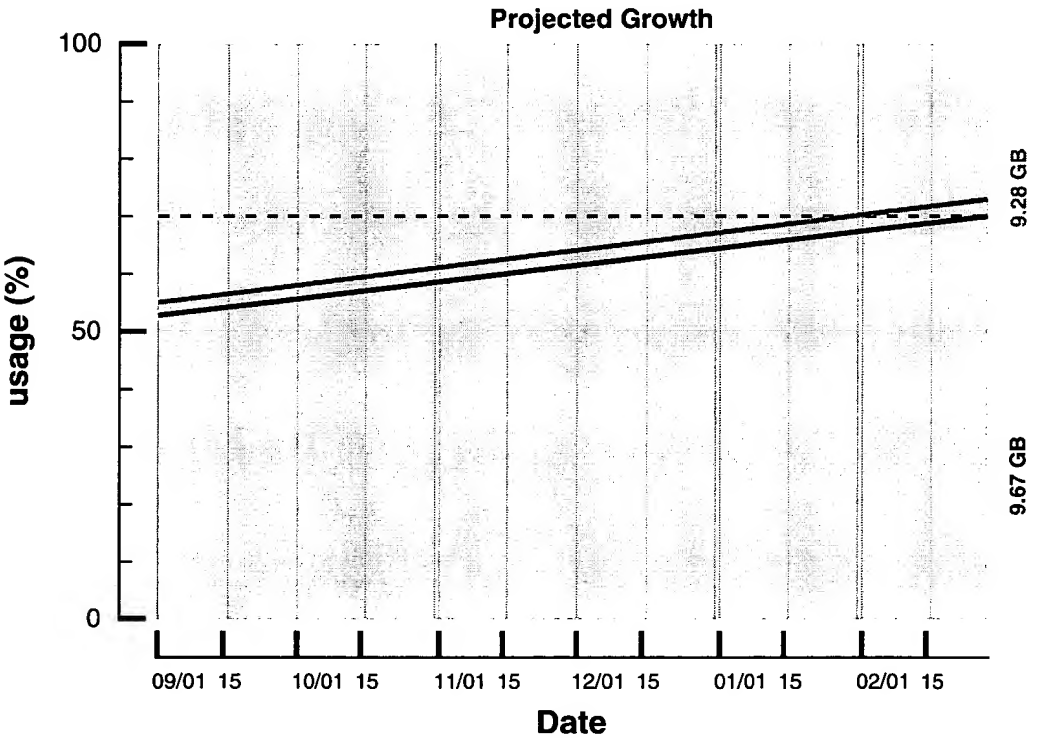

For this environment to operate satisfactorily in a future horizon of 180 days, it is necessary to add 1 disks, as shown in the graph below.

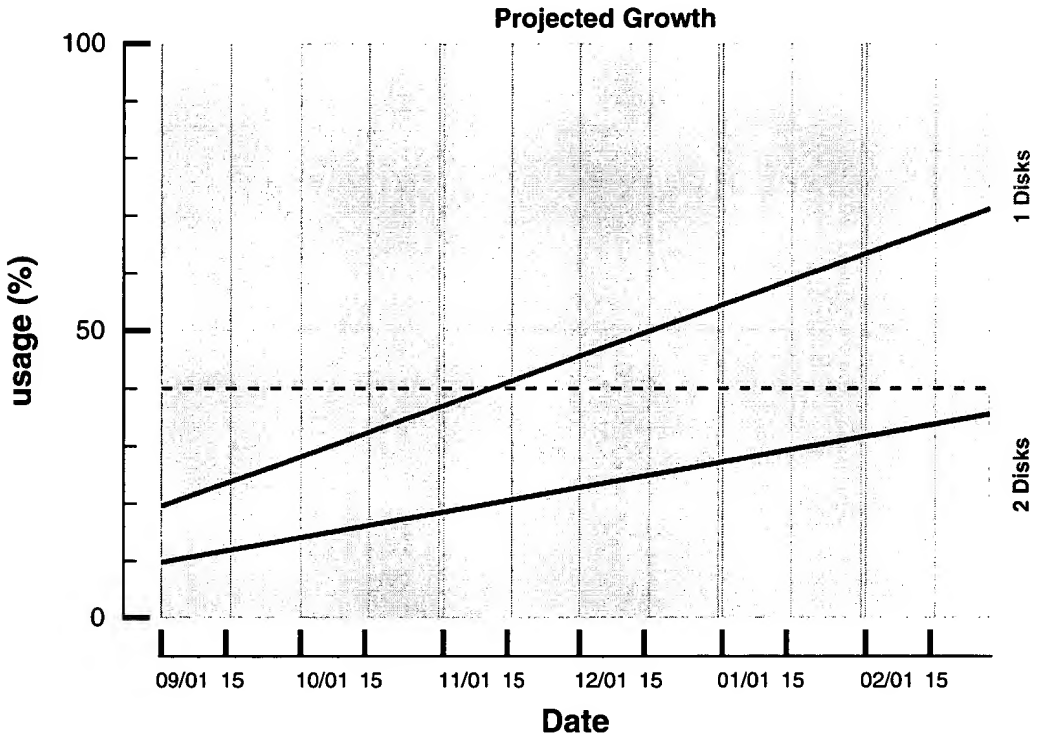

To understand a capacity planning report it is necessary to understand a few basic concepts. The idea is not to present a treaty, but to explain some fundamental aspects of capacity planning.

A computer installation may be represented by two compound systems, the user community and the computer system. The computer system is a hardware, software and transmission line complex destined to fulfill the processing and information needs of the users. These needs are communicated to the computer system through programs, data and commands produced by the users. This collection of programs, data and commands is called workload.

The computer system has a limited performance that may be quantified with measures such as: usage rate, response time, processing rate, availability index, etc. The performance of a computer system depends upon the interaction between the workload and the resources of the system. A new concept arises, the system's Capacity. This concept may be defined as the workload that a given computer system can process without exceeding the performance limits imposed by the installation.

When the workload surpasses the established limits, the system is said to have "exceeded the saturation limit", or the system is "saturated". Usually, from this moment on the response times to user requests become too slow or present an erratic behaviour.

The basic purpose of Capacity Planning is to provide, in due time, the necessary reports for rendering efficient Information Technology services .

There are different capacity planning techniques, depending basically on three factors: complexity, precision and cost. These techniques are the following, in order of growing complexity, precision and cost:

-empirical rules

-linear analysis

-analytic models

-simulation models

-"benchmarks"

Empirical rules are based on experience, knowledge, practice and feeling. It is cheap but very unreliable.

Linear Analysis, such as "Capacity Wizard", is based on the performance analysis for the current workload and the linear projection of the future behavior of this system from the present performance.

This method presents an excellent cost/benefit ratio, specially when applied in large scale in an automatic manner. One can have almost immediate capacity evaluations for hundreds of servers and workstations, with a minimum human intervention. These evaluations will permit rapid visualization of future bottlenecks in the system, permitting preventive measures, not mere reactions.

An additional advantage of Capacity Wizard is providing a constant and continued evaluation, before and after configuration changes. This means that, after an upgrade or improvement in the system, a new report will indicate if the results of the change were as good as predicted in the previous report (due to the great quantity of variables and unexpected occurrences involved, Capacity Planning is not an exact science).

The analytic model technique is based on a set of mathematical equations that

represent the structure and functioning of a computer system during a given interval of time. Operational analysis and stochastic modeling are the techniques used.

Simulation is a numerical method describing the dynamic behavior of a system through time. These models consume a lot of CPU resources until they reach the intended results.

Benchmarking consists in the selection of a group of applications that represent as closely as possible the total workload involved. This group will be processed in a system as similar as possible to the system in question. This technique has many complications, such as the true representation of the programs and data masses and of the equipment used.

This tool makes the following analysis and projections:

-CPU

-Memory

-Disks- occupation

-Disks- usage

-Network Cards

The agent, installed in the analyzed system, collects data globally and per process for each of these resources.

The tool searches, in the collector agent database, the data used in the analysis. These data are validated and rearranged in the best manner.

The hourly consumptions, for each resource, are measured. The greatest recorded consumption is chosen as representative for that day. If the sample is long enough- default period is 3 months- the greatest consumption of the week will represent weekly consumption.

After these consumptions are defined, the tool will execute their linear regression, defining a line segment that approximates the defined daily or weekly consumption.

Once this segment is defined, one can extend it into the future and estimate when saturation of the resource will occur (the point where this line crosses the saturation level).

It may be possible to find that the resource is already saturated.

The tool accepts several definitions for saturation level. If the costumer does not want to define them, the tool will use the manufacturer's measurements for CPU and memory, and 50% capacity for disks and card networks.

If the user wants to define his own limits, he may do so. In doing so, he may know in advance what would happen if he upgraded his equipment.

Besides informing the moment of saturation, the tool informs the necessary alteration to prevent this excess.

Example: How much must we upgrade the present CPU so that it does not become saturated for the next 18 months?

Again, there is a default option (1 year), for all the resources.

The tool provides different types of graphs and charts, showing which applications consumed the most resources.

The following options influence the calculation of future projections and increments (equipment upgrades or alterations) :

For CPU and Memory:

1) Increment based on the present equipment (default).Note: if the equipment can no longer be upgraded, the user must choose the following option. 2) User-determined increment (user determines how many times he wants performance improved-2x,1.5x,etc.). In this case, the tool will not have to search any hardware databases.

For disks and card networks:

1) Pre-determined default increment. This increment assumes that an identical equipment was placed alongside the original one and the consumption was divided in half. 2) User-determined increments(Note: user may also define negative increments, that is, a downgrade of the equipment).

Interval of time before saturation:

1)Pre-determined (default). This assumes that the resource will not become saturated before one year starting on the last measurement. If the resource does become saturated before one year, the report informs the necessary upgrade to avoid this. 2)User-determined. User determines how many days must pass before the resource becomes saturated.